

On Oblivious PTAS's for Nash Equilibrium*

Constantinos Daskalakis[†]
EECS and CSAIL, MIT
costis@mit.edu

Christos Papadimitriou[‡]
Computer Science, U.C. Berkeley
christos@cs.berkeley.edu

January 12, 2013

Abstract

If a class of games is known to have a Nash equilibrium with probability values that are either zero or $\Omega(1)$ — and thus with support of bounded size — then obviously this equilibrium can be found exhaustively in polynomial time. Somewhat surprisingly, we show that there is a PTAS for the class of games whose equilibria are guaranteed to have *small* — $O(\frac{1}{n})$ — values, and therefore large — $\Omega(n)$ — supports. We also point out that there is a PTAS for games with sparse payoff matrices, a family for which the exact problem is known to be PPAD-complete [6]. Both algorithms are of a special kind that we call *oblivious*: The algorithm just samples a fixed distribution on pairs of mixed strategies, and the game is only used to determine whether the sampled strategies comprise an ϵ -Nash equilibrium; the answer is “yes” with inverse polynomial probability (in the second case, the algorithm is actually deterministic). These results bring about the question: *Is there an oblivious PTAS for finding a Nash equilibrium in general games?* We answer this question in the negative; our lower bound comes close to the quasi-polynomial upper bound of [18].

Another recent PTAS for *anonymous* games [13, 14, 7] is also oblivious in a weaker sense appropriate for this class of games (it samples from a fixed distribution on *unordered* collections of mixed strategies), but its running time is exponential in $\frac{1}{\epsilon}$. We prove that any oblivious PTAS for anonymous games with two strategies and three player types must have $\frac{1}{\epsilon^\alpha}$ in the exponent of the running time for some $\alpha \geq \frac{1}{3}$, rendering the algorithm in [7] (which works with any bounded number of player types) essentially optimal within oblivious algorithms. In contrast, we devise a $\text{poly}(n) \cdot (1/\epsilon)^{O(\log^2(1/\epsilon))}$ non-oblivious PTAS for anonymous games with two strategies and any bounded number of player types. The key idea of our algorithm is to search not over unordered sets of mixed strategies, but over a carefully crafted set of collections of the first $O(\log \frac{1}{\epsilon})$ moments of the distribution of the number of players playing strategy 1 at equilibrium. The algorithm works because of a probabilistic result of more general interest that we prove: the total variation distance between two sums of independent indicator random variables decreases exponentially with the number of moments of the two sums that are equal, independent of the number of indicators.

A byproduct of our algorithm is establishing the existence of a sparse (and efficiently computable) ϵ -cover of the set of all possible sums of n independent indicators, under the total variation distance. The size of the cover is $\text{poly}(n) \cdot (1/\epsilon)^{O(\log^2(1/\epsilon))}$.

*Extended version of the paper of the same title that appeared in *STOC*, 2009.

[†]This work was done while the author was a postdoctoral researcher at Microsoft Research, New England.

[‡]Supported by NSF grant CCF - 0635319, a gift from Yahoo! Research, and a MICRO grant.

1 Introduction

Is there a polynomial time approximation scheme (PTAS) for computing approximate Nash equilibria? This has emerged, in the wake of the intractability results for Nash equilibria [9, 5], as the most central question in equilibrium computation. Over the past three years there has been relatively slow progress towards smaller ϵ 's [10, 11, 4, 21] — so slow that it is hard to believe that a PTAS is around the corner. On the other hand, [13, 14] provide a PTAS for the important special case of anonymous games with a bounded number of strategies (those for which the utilities of the players, although different, depend on the *number* of players playing each strategy, not on the *identities* of the players that do). This PTAS proceeds by discretizing the probabilities in the mixed strategies of the players to multiples of $\frac{1}{k}$, for appropriate integer k , and works even in the generalization in which the players are divided into a bounded number of *types*, and utilities depend on how many players *of each type* choose each strategy.

In this paper we report significant progress on this important algorithmic problem; we present several new algorithms, but also the first nontrivial lower bounds. We start by pointing out two new interesting classes of bimatrix games¹ that have PTAS's:

- It was shown in [6] that computing a Nash equilibrium for the special case of *sparse games*, that is, games whose payoff matrices have a bounded number of non-zero entries in each row and column [6], is PPAD-complete. We point out that there is a trivial PTAS — in particular, the pair of uniform mixed strategies works. This is interesting in that this is the first PPAD-complete case of a problem that is so approximable.
- We also give a randomized PTAS for *small probability games*, that is, games that are guaranteed to have Nash equilibria with small $O(\frac{1}{n})$ nonzero probability values, and thus with linear support (Theorem 3).

It is quite surprising that games with small probability values (our second special case above) are easy, since games with *large* (bounded from below by a constant) probability values are also easy. What probability values are difficult then? Our next result, a lower bound, seems to suggest *inverse logarithmic* probability values are hard (compare with the quasi-PTAS of [18], whose equilibria have roughly logarithmic support).

To explain our negative result, we first note that both PTAS's outlined above (as well as those for anonymous games discussed later) are *oblivious*. This means that they have access to a fixed set of pairs of mixed strategies, in the generic case by sampling, and they look at the game only to determine whether a sampled pair of strategies constitutes an approximate Nash equilibrium. Is there an oblivious PTAS for the general Nash equilibrium? The answer here is a definite “no” — in fact, we proved our negative result after we had been working for some time on developing such an algorithm. . . We show that any oblivious algorithm must sample at least $\Omega(n^{(.8-34\epsilon)\log_2 n})$ pairs in expectation (Theorem 4). For comparison, [18]'s algorithm takes time $n^{O(\log n/\epsilon^2)}$.

Another important class of games for which a PTAS was recently discovered is that of *anonymous games with two strategies* [13, 7]. These are multi-player games, in which the utility of each player depends on the strategy (0 or 1) played, and the *number* of other players playing strategy 1 (*not* their identities). In fact, the PTAS works even in the more sophisticated case in which the players are partitioned into types, and the utilities depend on the number of players *of each type* playing strategy 1. The PTAS in [13], and the relatively more efficient one in [7], both have running times that are exponential in $\frac{1}{\epsilon}$. They are also oblivious in a sense appropriate for anonymous

¹These results can be extended to any bounded number of players, but in what follows we only discuss the two-player, or bimatrix, case.

games, in that they work by sampling an *unordered* set of n mixed strategies, where n is the number of players, and they only look at the game in order to determine if there is an assignment of these strategies to the players that results in an approximate equilibrium. We prove that any oblivious approximation algorithm, for anonymous games with two strategies and three player types, must sample an exponential—in $\frac{1}{\epsilon}$ —sized collection of unordered sets of mixed strategies — and so our PTAS’s are near-optimal.

Finally, we circumvent this negative result to develop a non-oblivious PTAS which finds an ϵ -approximate Nash equilibrium in anonymous games with two strategies and any bounded number of player types in time $\text{poly}(n) \left(\frac{1}{\epsilon}\right)^{O(\log^2 \frac{1}{\epsilon})}$. This algorithm is based (in addition to many other insights and techniques for anonymous games) on a new result in applied probability which is, we believe, interesting in its own right: Suppose that you have two sums of n independent Bernoulli random variables, which have the same first moment, the same second moment, and so on up to moment d . Then the distributions of the two sums have variational distance that vanishes exponentially fast with d , regardless of n . To turn this theorem into an algorithm, we discretize the mixed strategies of the players using techniques from [7] and, in the range of parameters where the algorithm of [7] breaks down, we iterate over all possible values of the first $\log(1/\epsilon)$ moments of the players’ aggregate behavior; we then try to identify, via an involved dynamic programming scheme, mixed strategies, implementing the given moments, which correspond to approximate Nash equilibria of the game. Our approximation guarantee for sums of independent indicators is rather strong, especially when the number of indicators is small, a regime where Berry-Esséen type bounds provide weaker guarantees and result in slower algorithms [13, 7]. It is quite intriguing that a quasi-polynomial time bound, such as the one we provide in this paper, shows up again in the analysis of algorithms for approximate Nash equilibria (cf. [18]).

As a byproduct of our results we establish the existence of a sparse (and efficiently computable) ϵ -cover of the set of all possible sums of n independent indicators, under the total variation distance. The size of our cover is $\text{poly}(n) \cdot (1/\epsilon)^{O(\log^2(1/\epsilon))}$. We discuss this result in Section 6.

1.1 Preliminaries

A two-player, or *bimatrix*, game \mathcal{G} is played by two players, the *row player* and the *column player*. Each player has a set of n *pure strategies*, which without loss of generality we assume to be the set $[n] := \{1, \dots, n\}$ for both players. The game is described then by two *payoff matrices* R, C corresponding to the row and column player respectively, so that, if the row player chooses strategy i and the column player strategy j , the row player gets payoff R_{ij} and the column player C_{ij} . As it is customary in the literature of approximate Nash equilibria, we assume that the matrices are normalized, in that their entries are in $[-1, 1]$.

The players can play *mixed strategies*, that is, probability distributions over their pure strategies which are represented by probability vectors $x \in \mathbb{R}_+^n$, $|x|_1 = 1$. If the row player chooses mixed strategy x and the column player mixed strategy y , then the row player gets expected payoff $x^T R y$ and the column player expected payoff $x^T C y$. A pair of mixed strategies (x, y) is a *Nash equilibrium* of the game $\mathcal{G} = (R, C)$ iff x maximizes $x'^T R y$ among all probability vectors x' and, simultaneously, y maximizes $x^T C y'$ among all y' . It is an ϵ -approximate Nash equilibrium iff $x^T R y \geq x'^T R y - \epsilon$, for all x' , and, simultaneously, $x^T C y \geq x^T C y' - \epsilon$, for all y' . In this paper, we will use the stronger notion of ϵ -approximately well-supported Nash equilibrium, or simply ϵ -Nash equilibrium. This is any pair of strategies (x, y) such that, for all i with $x_i > 0$, $e_i^T R y \geq e_{i'}^T R y - \epsilon$, for all i' , and similarly for y . That is, every strategy i in the support of x guarantees expected payoff $e_i^T R y$ which is within ϵ from the optimal response to y , and similarly every strategy in the support of y is within ϵ from the optimal response to x .

An *anonymous game* is in a sense the dimensional dual of a bimatrix game: There are n players, each of which has two strategies, 0 and 1. For each player i there is a *utility function* u_i mapping $\{0, 1\} \times [n - 1]$ to $[-1, 1]$. Intuitively, $u_i(s, k)$ is the utility of player i when s/he plays strategy $s \in \{0, 1\}$, and $k \leq n - 1$ of the remaining players play strategy 1, while $n - k - 1$ play strategy 0. In other words, the utility of each player depends, in a player-specific way, on the strategy played by the player and the number of other players who play strategy 1 — but *not* their identities. The notions of Nash equilibrium and ϵ -Nash equilibrium are extended in the natural way to this setting. Briefly, a mixed strategy for the i -th player is a function $x_i : \{0, 1\} \rightarrow [0, 1]$ such that $x_i(0) + x_i(1) = 1$. A set of mixed strategies x_1, \dots, x_n is then an ϵ -Nash equilibrium if, for every player i and every $s \in \{0, 1\}$ with $x_i(s) > 0$: $\mathcal{E}_{x_1, \dots, x_n} u_i(s, k) \geq \mathcal{E}_{x_1, \dots, x_n} u_i(1 - s, k) - \epsilon$, where for the purposes of the expectation k is drawn from $\{0, \dots, n - 1\}$ by flipping $n - 1$ coins according to the distributions $x_j, j \neq i$, and counting the number of 1's.

A more sophisticated kind of anonymous games divides the players into t types, so that the utility of player i depends on the strategy played by him/her, and the number of players of *each type* who play strategy 1.

2 PTAS for Two Special Cases

2.1 Small Games

We say that a class of bimatrix games is *small* if the sum of all entries of the R and C matrices is $o(n^2)$. One such class of games are the k -sparse games [6], in which every row and column of both R and C have at most k non-zero entries. The following result by Chen, Deng and Teng shows that finding an exact Nash equilibrium remains hard for k -sparse games:

Theorem 1 ([6]). *Finding a $O(n^{-6})$ -Nash equilibrium in 10-sparse normalized bimatrix games with n strategies per player is a PPAD-complete problem.*

In contrast, it is easy to see that there is a PTAS for this class:

Theorem 2. *For any k , there is a PTAS for the Nash equilibrium problem in k -sparse bimatrix games.*

Our original proof of this theorem consisted in showing that there always exists an ϵ -Nash equilibrium in which both players of the game use in their mixed strategies probabilities that are integer multiples of $\epsilon/2k$. Hence, we can efficiently enumerate over all possible pairs of mixed strategies of this form, as long as k is fixed. Shang-Hua Teng pointed out to us a much simpler algorithm: The pair of uniform mixed strategies is always an ϵ -Nash equilibrium in a sparse game! The difference with our algorithm is this: the uniform equilibrium gives to both players payoff of at most k/n ; our algorithm can be used instead to approximate the payoffs of the players in the Nash equilibrium with the optimal social welfare (or more generally the Nash equilibrium that optimizes some other smooth function of the players' payoffs).

2.2 Small Probability Games

If all games in a class are guaranteed to have a Nash equilibrium (x, y) , where the nonzero probability values in both x and y are larger than some constant $\delta > 0$, then it is trivial to find this equilibrium in time $n^{O(\frac{1}{\delta})}$ by exhaustive search over all possible supports and linear programming. But, what if a class of games is known to have *small*, say $O(\frac{1}{n})$, probability values? Clearly, exhaustive search over supports is not efficient anymore, since those have now linear size. Surprisingly, we show that any such class of games has a (randomized) PTAS, by exploiting the technique of [18].

Definition 1 (Small Probability Games). *A bimatrix game $\mathcal{G} = (R, C)$ is of δ -small probabilities, for some constant $\delta \in (0, 1]$, if it has a Nash equilibrium (x, y) such that all the entries of x and y are at most $\frac{1}{\delta n}$.*

Remark 1. *Observe that a game of δ -small probabilities has an equilibrium (x, y) , in which both x and y have support of size at least δn . Moreover, there exists a subset of size at least $\frac{\delta n}{2}$ from the support of x , all the strategies of which have probability at least $\frac{1}{2n}$, and similarly for y ; that is the probability mass of the distributions x and y spreads non-trivially over a subset of size $\Omega(n)$ of the strategies. Hence, small probability games comprise a subclass of linear support games, games with an equilibrium (x, y) in which both x and y assign to a constant fraction αn of the strategies probability at least $1/\beta n$, for some constants α and β . However, this broader class of games is essentially as hard as the general: take any Nash equilibrium (x, y) of a normalized game and define the pair (x', y') , where $x' := (1 - \frac{\epsilon}{5}) \cdot x + \frac{\epsilon}{5} \cdot (\frac{1}{n}, \frac{1}{n}, \dots, \frac{1}{n})$ and similarly for y' . It is not hard to see that the new pair is an ϵ -Nash equilibrium; still, regardless of what (x, y) is, both x' and y' assign to αn strategies probability at least $1/\beta n$, for an appropriate selection of α and β .*

Theorem 3. *For any $\delta \in (0, 1]$, there is a randomized PTAS for normalized bimatrix games of δ -small probabilities.*

Proof. We show first the following (stronger in terms of the type of approximation) variant of the theorem of Lipton, Markakis and Mehta [18]. The proof is provided in Appendix A.1.

Lemma 1. *Let $\mathcal{G} = (R, C)$ be a normalized bimatrix game and (x, y) a Nash equilibrium of \mathcal{G} . Let \mathcal{X} be the distribution formed by taking $t = \lceil 16 \log n / \epsilon^2 \rceil$ independent random samples from x and defining the uniform distribution on the samples, and similarly define \mathcal{Y} by taking samples from y . Then with probability at least $1 - \frac{4}{n}$ the following are satisfied*

1. *the pair $(\mathcal{X}, \mathcal{Y})$ is an ϵ -Nash equilibrium of \mathcal{G} ;*
2. *$|e_i^T R \mathcal{Y} - e_i^T R y| \leq \epsilon/2$, for all $i \in [n]$;*
3. *$|\mathcal{X}^T C e_j - x^T C e_j| \leq \epsilon/2$, for all $j \in [n]$;*

Suppose now that we are given a normalized bimatrix game $\mathcal{G} = (R, C)$ of δ -small probabilities, and let (x, y) be an equilibrium of \mathcal{G} in which $x_i \leq \frac{1}{\delta n}$, for all $i \in [n]$, and similarly for y . Lemma 1 asserts that, if a multiset² A of size $t = \lceil 16 \log n / \epsilon^2 \rceil$ is formed by taking t independent samples from x and, similarly, a multiset B is formed by taking samples from y , then $(\mathcal{X}, \mathcal{Y})$, where \mathcal{X} is the uniform distribution over A and \mathcal{Y} the uniform distribution over B , is an ϵ -Nash equilibrium with probability at least $1 - 4/n$. Of course, we do not know (x, y) so we cannot do the sampling procedure described above. Instead we are going to take a uniformly random multiset A' and a uniformly random multiset B' and form the uniform distributions $\mathcal{X}', \mathcal{Y}'$ over A' and B' ; we will argue that there is an inverse polynomial chance that $(\mathcal{X}', \mathcal{Y}')$ is actually an ϵ -Nash equilibrium.

For this we define the set \mathcal{A} of *good multisets for the row player* as

$$\mathcal{A} := \left\{ A \left| \begin{array}{l} A \text{ is a multiset, } A \subseteq [n], |A| = t, \text{ the} \\ \text{uniform distribution } \mathcal{X} \text{ over } A \text{ satisfies} \\ \text{Assertion 3 of Lemma 1} \end{array} \right. \right\},$$

²For our discussion, a multiset of size t is an *ordered* collection $\langle i_1, i_2, \dots, i_t \rangle$ of t elements from some universe (repetitions are allowed).

and, similarly, the set \mathcal{B} of *good multisets of the column player* as

$$\mathcal{B} := \left\{ B \left| \begin{array}{l} B \text{ is a multiset, } B \subseteq [n], |B| = t, \text{ the} \\ \text{uniform distribution } \mathcal{Y} \text{ over } B \text{ satisfies} \\ \text{Assertion 2 of Lemma 1} \end{array} \right. \right\}.$$

The reason for defining \mathcal{A} and \mathcal{B} in this way is that, given two multisets $A \in \mathcal{A}$, $B \in \mathcal{B}$, the uniform distributions \mathcal{X} over A and \mathcal{Y} over B comprise an ϵ -Nash equilibrium (see the proof of Lemma 1 for a justification).

What remains to show is that, with inverse polynomial probability, a random multiset A' belongs to \mathcal{A} and a random multiset B' belongs to \mathcal{B} . To show this we lower bound the cardinalities of the sets \mathcal{A} and \mathcal{B} via the following claim, proven in Appendix A.1. We argue that the subset of \mathcal{A} containing elements that could arise by sampling x is large: indeed, with probability at least $1 - \frac{4}{n}$, a multiset sampled from x belongs to \mathcal{A} and, moreover, each individual multiset has small probability of being sampled, since x spreads the probability mass approximately evenly on its support.

Claim 1. *The sets \mathcal{A} and \mathcal{B} satisfy*

$$|\mathcal{A}| \geq \left(1 - \frac{4}{n}\right) (\delta n)^t \text{ and } |\mathcal{B}| \geq \left(1 - \frac{4}{n}\right) (\delta n)^t.$$

Given Claim 1, we can show Claim 2; the proof is given in Appendix A.1. Equation (1) implies that the algorithm that samples two uniformly random multisets A' , B' and forms the uniform probability distributions \mathcal{X}' and \mathcal{Y}' over A' and B' respectively, succeeds in finding an ϵ -Nash equilibrium with probability inverse polynomial in n . This completes the proof of Theorem 3.

Claim 2. *If $\mathcal{X}', \mathcal{Y}'$ are the uniform distributions over random multisets A' and B' then*

$$\begin{aligned} \Pr [(\mathcal{X}', \mathcal{Y}') \text{ is an } \epsilon\text{-Nash equilibrium}] \\ = \Omega \left(\delta^2 \cdot n^{-32 \log(1/\delta)/\epsilon^2} \right). \end{aligned} \tag{1}$$

□

3 A Lower Bound for Bimatrix Games

The two PTAS's presented in the previous section are *oblivious*. An oblivious algorithm looks at the game only to check if the various pairs of mixed strategies it has come up with (by enumeration or, more generally, by random sampling) are ϵ -approximate, and is guaranteed to come up with one that is with probability at least inverse polynomial in the game description. More formally, an oblivious algorithm for the Nash equilibrium problem is a distribution over pairs of mixed strategies, indexed by the game size n and the desired approximation ϵ . It is a PTAS if for any game the probability that a pair drawn from the distribution is an ϵ -Nash equilibrium is inversely polynomial in n . Notice that, since we are about to prove lower bounds, we are opting for the generality of randomized oblivious algorithms—a deterministic algorithm that enumerates over a fixed set of pairs of mixed strategies can be seen as a (randomized) oblivious algorithm by considering the uniform distribution over the set it enumerates over.

The rather surprising simplicity and success of these algorithms (as well as their cousins for anonymous games, see the next section) raises the question: Is there an oblivious PTAS for the general Nash equilibrium problem? We show that the answer is negative.

Theorem 4. *There is no oblivious PTAS for the Nash equilibrium in bimatrix games.*

Proof. We construct a super-polynomially large family of $n \times n$ games with the property that every two games in the family do not share an ϵ -Nash equilibrium. This quickly leads to the proof of the theorem.

Our construction is based on a construction by Althöfer [1], except that we need to pay more attention to ensure that the ϵ -Nash equilibria of the games we construct are “close” to the exact Nash equilibria. For ℓ even and $n = \binom{\ell}{\ell/2}$, we define a family of $n \times n$ two-player games $\mathcal{G}_S = (R_S, C_S)$, indexed by all subsets $S \subseteq [n]$ with $|S| = \ell$. Letting $\{S_1, S_2, \dots, S_n\}$ be the set of all subsets of S with cardinality $\ell/2$, we imagine that column j of the game \mathcal{G}_S corresponds to subset S_j . Then, for every j , we fill column j of the payoff matrices R_S and C_S as follows:

- for all $i \notin S$, $R_{S,ij} = -1$ and $C_{S,ij} = 1$;
- for all $i \in S_j$, $R_{S,ij} = 1$ and $C_{S,ij} = 0$; and
- for all $i \in S \setminus S_j$, $R_{S,ij} = 0$ and $C_{S,ij} = 1$.

In other words, our construction has two components: In the first component (defined by the rows labeled with the elements of S), the game is 1-sum, whereas in the second (corresponding to the complement of S) the game is 0-sum with the row player always getting payoff of -1 and the column player always getting payoff of 1 . The payoffs of the first component are more balanced in the following way: as we said, every column corresponds to a subset S_1, \dots, S_n of S of cardinality $\ell/2$; if the column player chooses column j , then the row player gets 1 for choosing a row corresponding to an element of S_j and 0 for a row corresponding to an element in $S \setminus S_j$. See Figure 1 of Appendix A.2 for an illustration of R_S for the case $n = 6$, $\ell = 4$, $S = \{1, 2, 3, 4\}$.

Lemma 2 provides the following characterization of the approximate equilibria of the game \mathcal{G}_S : in any ϵ -Nash equilibrium (x, y) , the strategy x of the row player must have ℓ_1 distance at most 8ϵ from u_S —the uniform distribution over the set S . That is, in all approximate equilibria of the game the strategy of the row player must be close to the uniform distribution over S . Formally,

Lemma 2. *Let $\epsilon < 1$. If (x, y) is an ϵ -Nash equilibrium of the game \mathcal{G}_S , where x is the mixed strategy of the row player and y that of the column player, then the following properties are satisfied by x :*

1. $x_i = 0$, for all $i \notin S$;
2. $\ell_1(x, u_S) \leq 8\epsilon$, where u_S is the uniform distribution over S , and $\ell_1(x, u_S)$ represents the ℓ_1 distance between distributions x , u_S .

The proof of the first assertion is straightforward: the row player will not assign any probability mass to the rows which give her -1 , since any row in S will guarantee her at least 0 . Since all the activity happens then in the first component of the game, which is 1-sum, an averaging argument implies that both players’ payoff is about $1/2$ at equilibrium. Observe further that, for a given mixed strategy of the row player, the strategy of the column player that guarantees her the highest payoff is the subset containing the $\ell/2$ elements of S to which the row player assigns the smallest probability mass. Hence, if the probability distribution of the row player were far from uniform, then, contrary to what we argued, the corresponding payoff for the column player would be larger than $1/2$ —this is established via a delicate geometric argument for ℓ_1 distances of probability measures. See Lemma 3 in Appendix A.2.

Suppose now that there is an oblivious PTAS for the Nash equilibrium, that is, a distribution D over pairs of mixed strategies such that, for any game \mathcal{G}_S , the probability that an ϵ -approximate

Nash equilibrium is sampled is inverse polynomial in n . Let us consider the probability distribution D_R induced by D on the mixed strategies of the row player and denote by B_S the ℓ_1 ball of radius 8ϵ around u_S . Lemma 2 implies that D_R should be assigning probability mass at least inverse polynomial in n to each ball B_S , $S \subseteq [n]$, $|S| = \ell$. This is impossible, since by the following claim there is a super-polynomial number of disjoint such balls. The proof of the claim is via a counting argument. See Appendix A.2.

Claim 3. *There is a family of $\Omega(n^{(.8-34\epsilon)\log_2 n})$ disjoint such balls.*

□

The proof of Theorem 4 implies in particular that any oblivious PTAS for general two-player games requires expected running time of at least $\Omega(n^{(.8-34\epsilon)\log_2 n})$. Compare this bound to the $n^{O(\log n/\epsilon^2)}$ upper bound obtained by Lipton, Markakis and Mehta [18].

4 A Lower Bound for Anonymous Games

Recall the definition of anonymous games from Section 1.1. In [13] we give a PTAS for anonymous games with two strategies, running in time $n^{O((1/\epsilon)^2)}$, and in [7] a more efficient one, with running time $n^{O(1)} \cdot (1/\epsilon)^{O((1/\epsilon)^2)}$. (In [14] we also give a much more sophisticated PTAS for anonymous games with more than two strategies). All these PTAS's have $\frac{1}{\epsilon}$ in the exponent of the running time, and they work even if there is a fixed number t of types (in which case t multiplies the exponent). Furthermore, it turns out that all of these algorithms are *oblivious*, in a specialized sense appropriate for anonymous games defined next.

An *oblivious ϵ -approximation algorithm for anonymous games with n players* is defined in terms of a distribution, indexed by n , on *unordered n -tuples of mixed strategies*.³ The algorithm samples from this distribution, and for each $\{p_1, \dots, p_n\}$ sampled, it determines whether there is an assignment of these probabilities to the n players such that the resulting strategy profile (with each player playing strategy 1 with the assigned probability) is an ϵ -approximate Nash equilibrium; this latter test can be carried out by max-flow techniques, see, e.g., [13]. The expected running time of this approximation algorithm is then the inverse of its probability of success.

We show the following result, implying that any oblivious ϵ -approximation algorithm for anonymous games whose expected running time is polynomial in the number of players must have expected running time exponential in $(\frac{1}{\epsilon})^{1/3}$. Hence, our PTAS from [7] is essentially optimal among oblivious PTAS's.

Theorem 5. *For any constant $c \geq 0$, $\epsilon < 1$, no oblivious ϵ -approximation algorithm for anonymous games with 2 strategies and 3 player types has probability of success larger than $n^{-c} \cdot 2^{-\Omega(1/\epsilon^{1/3})}$.*

We only sketch the proof next and postpone the details for Appendix B. We first show the following (see Theorem 10 in Appendix B.1): given any ordered n -tuple (p_1, \dots, p_n) of probabilities, we can construct a normalized anonymous game with n players of type A, and two more players of their own type, such that in any ϵ -Nash equilibrium of the game the i -th player of type A plays strategy 1 with probability very close (depending on ϵ and n) to p_i . To obtain this game, we need to understand how to exploit the difference in the payoff functions of the players of type A to enforce different behaviors at equilibrium, despite the fact that in all other aspects of the game the players of group A are indistinguishable.

³For ordered sets of mixed strategies, the lower bound we are about to show becomes trivial and uninteresting.

The construction is based on the following idea: For all i , let us denote by $\mu_{-i} := \sum_{j \neq i} p_j$ the *target* expected number of type-A players different than i who play strategy 1; and let us give this payoff to player i if she plays strategy 0, regardless of what the other players are doing. If i chooses 1, we give her payoff t where t is the number of players different than i who play 1. By setting the payoffs in this way we ensure that (p_1, \dots, p_n) is in fact an equilibrium, since for every player the payoff she gets from strategy 0 matches the expected payoff she gets from strategy 1. However, enforcing that (p_1, \dots, p_n) is also the unique equilibrium is a more challenging task; and to do this we need to include two other players of their own type: we use these players to ensure that the sum of the mixed strategies of the players of type A matches $\sum p_i$ at equilibrium, so that a player i deviating from her prescribed strategy p_i is pushed back towards p_i . We show how this can be done in Appendix B.1. We also provide guarantees for the ϵ -Nash equilibria of the resulting game.

The construction outlined above enables us to define a family of $2^{\Omega(1/\epsilon^{1/3})}$ anonymous games with the property that no two games in the family share an ϵ -Nash equilibrium, even as an unordered set of mixed strategies (Claims 7 and 8). Then, by an averaging argument, we can deduce that for any oblivious algorithm there is a game in the ensemble for which the probability of success is at most $2^{-\Omega(1/\epsilon^{1/3})}$. It is important for our construction to work that the anonymous game defined for a given collection of probability values (p_1, p_2, \dots, p_n) does not deviate too much from the prescribed set of mixed strategies p_1, p_2, \dots, p_n even in an ϵ -Nash equilibrium. The bound of $2^{\Omega(1/\epsilon^{1/3})}$ emerges from a quantification of this deviation as a function of ϵ . The proof of Theorem 5 is given in Appendix B.2.

Remark 2. *We can show an equivalent of Theorem 5 for oblivious ϵ -approximation algorithms for anonymous games with 2-player types and 3 strategies per player. The details are omitted.*

5 A Quasi-polynomial PTAS

We circumvent the lower bound of the previous section by providing a PTAS for anonymous games with running time polynomial in the number of players n times a factor of $(\frac{1}{\epsilon})^{O(\log^2 \frac{1}{\epsilon})}$. The PTAS is, of course, non-oblivious, and in fact in the following interesting way: Instead of enumerating a fixed set of unordered collections of probability values, we enumerate a fixed set of $\log(1/\epsilon)$ -tuples, *representing the first $\log(1/\epsilon)$ moments of these probability values*. We can think of these moments as *more succinct aggregates* of mixed strategy profiles than the unordered collections of strategies considered in [7, 13], since several of these collections may share the same moments. To put the idea into context, let us recall the following theorem.

Theorem 6 ([7, 8]). *For every $\epsilon > 0$, there exists some integer $k = O(1/\epsilon)$ such that for any n -player 2-strategy anonymous game with payoffs in $[-1, 1]$ there is an ϵ -Nash equilibrium such that*

1. *either at most k^3 players randomize, and their mixed strategies are integer multiples of $1/k^2$;*
2. *or all players who randomize use the same mixed strategy, and this strategy is an integer multiple of $\frac{1}{kn}$.*

This structural result can be turned into an oblivious PTAS using max-flow arguments (see Appendix D.1 for details). At its heart the proof of the theorem relies on the following intuitive fact about sums of indicator random variables: *If two sums of independent indicators have close means and variances, then their total variation distance should be small.* The way this fact becomes

relevant to anonymous games is that, if there are 2 strategies per player, then the mixed strategy of a player can be described by an indicator random variable; and as it turns out, if we replace one set of indicators by another, the change in payoff that every player will experience is bounded by the total variation distance between the sum of the indicators before and after the change. Nevertheless, the bound obtained by approximating the first two moments of the sum of the Nash equilibrium strategies is weak, and the space of unordered sets of probability values that we need to search over becomes exponential in $1/\epsilon$.

To obtain an exponential pruning of the search space, we turn to higher moments of the Nash equilibrium. We show the following theorem, which provides a rather strong quantification of how the total variation distance between two sums of indicators depends on the number of their first moments that are equal.

Theorem 7. *Let $\mathcal{P} := (p_i)_{i=1}^n \in (0, 1/2]^n$ and $\mathcal{Q} := (q_i)_{i=1}^n \in (0, 1/2]^n$ be two collections of probability values in $(0, 1/2]$. Let also $\mathcal{X} := (X_i)_{i=1}^n$ and $\mathcal{Y} := (Y_i)_{i=1}^n$ be two collections of independent indicators with $\mathcal{E}[X_i] = p_i$ and $\mathcal{E}[Y_i] = q_i$, for all $i \in [n]$. If for some $d \in [n]$ the following condition is satisfied:*

$$(C_d) : \quad \sum_{i=1}^n p_i^\ell = \sum_{i=1}^n q_i^\ell, \quad \text{for all } \ell = 1, \dots, d,$$

$$\text{then } \left\| \sum_i X_i ; \sum_i Y_i \right\| \leq 20(d+1)^{1/4} 2^{-(d+1)/2}. \quad (2)$$

Condition (C_d) considers the power sums of the expectations of the indicator random variables. Using the theory of symmetric polynomials we can show that (C_d) is equivalent to the following condition on the moments of the sums of the indicators (for the proof see Proposition 1 in Appendix C):

$$(V_d) : \quad \mathcal{E} \left[\left(\sum_{i=1}^n X_i \right)^\ell \right] = \mathcal{E} \left[\left(\sum_{i=1}^n Y_i \right)^\ell \right], \quad \text{for all } \ell \in [d].$$

Theorem 7 provides then the following strong approximation guarantee for sums of indicator random variables, that we think should be important in other settings. Our result is related to the classical moment method in probability theory [15], but to our knowledge it is novel and significantly stronger than known results:

If two sums of independent indicators with expectations bounded by $1/2$ have equal first d moments, then their total variation distance is $2^{-\Omega(d)}$.

It is important to note that our bound in (2) does not rely on summing up a large number of indicators n . This is quite critical since the previous techniques break down for small n 's—for large n 's Berry-Esséen type bounds are sufficient to obtain strong guarantees (this is the heart of the probabilistic results used in [13, 7]).

The proof of Theorem 7 (and its complement for probability values in $[1/2, 1)$) is given in Appendix C. It proceeds by expressing the distribution of the sum of n indicators, with expectations p_1, \dots, p_n , as a weighted sum of the binomial distribution $\mathcal{B}(n, p)$ (with $p = \bar{p} = \sum p_i/n$) and its n derivatives with respect to p , at the value $p = \bar{p}$ (these derivatives correspond to finite signed measures). It turns out that the coefficients of the first d terms of this expansion are symmetric polynomials with respect to the probability values p_1, \dots, p_n , of degree at most d ; hence, from the theory of symmetric polynomials, each of these coefficients can be written as a function of the power-sum symmetric polynomials $\sum_i p_i^\ell$, $\ell = 1, \dots, d$ (see, e.g., [22]). So, if two sums of indicators

satisfy Condition (C_d) , the first terms cancel, and the total variation distance of the two sums depends only on the other terms of the expansion (those corresponding to higher derivatives of the binomial distribution). The proof is concluded by showing that the joint contribution of these terms is inverse polynomial in $2^{\Omega(d)}$. (The details are given in Appendix C).

Our algorithm, shown below, exploits the strong approximation guarantee of Theorem 7 to improve upon the algorithm of [7] in the case where only $O(1/\epsilon^3)$ players mix at the ϵ -Nash equilibrium (this corresponds to case 1 in Theorem 6). The complementary case (case 2 in Theorem 6) can be treated easily in time polynomial in n and $1/\epsilon$ by exhaustive search and max-flow arguments (see Appendix D.1).

ALGORITHM MOMENT SEARCH

Input: An anonymous game \mathcal{G} , the desired approximation ϵ .

Output: An ϵ -Nash equilibrium of \mathcal{G} in which all probabilities are integer multiples of $\frac{1}{k^2}$, where $k = \lceil \frac{c}{\epsilon} \rceil$ and c is universal (independent of n) constant, determined by Theorem 6. For technical reasons that will be clear shortly, we choose a value for k that is by a factor of 2 larger than the value required by Theorem 6; this is the value k that guarantees an $\epsilon/2$ -Nash equilibrium in multiples of $1/k^2$. Finally, we assume that we have already performed the search corresponding to Case 2 of Theorem 6 for this value of k and we have not found an $\epsilon/2$ -Nash equilibrium. So there must exist an $\epsilon/2$ -Nash equilibrium in which at most k^3 players randomize in integer multiples of $1/k^2$.

1. Guess integers $t_0, t_1, t_s, t_b \leq n$, $t_s + t_b \leq k^3$, where t_0 players will play pure strategy 0, t_1 will play pure strategy 1, t_s will mix with probability $\leq \frac{1}{2}$, and $t_b = n - t_0 - t_1 - t_s$ will mix with probability $> \frac{1}{2}$. (Note that we have to handle low and high probabilities separately, because Theorem 7 only applies to indicators with expectations in $(0, 1/2]$; we handle indicators with expectations in $(1/2, 1)$ by taking their complement and employing Theorem 7—see Corollary 1.)
2. For $d = \lceil 3 \log_2(320/\epsilon) \rceil$, guess $\mu_1, \mu_2, \dots, \mu_d, \mu'_1, \mu'_2, \dots, \mu'_d$, where, for all $\ell \in [d]$:

$$\mu_\ell \in \left\{ j \left(\frac{1}{k^2} \right)^\ell : t_s \leq j \leq t_s \left(\frac{k^2}{2} \right)^\ell \right\},$$

and

$$\mu'_\ell \in \left\{ j \left(\frac{1}{k^2} \right)^\ell : t_b \left(\frac{k^2}{2} + 1 \right)^\ell \leq j \leq t_b (k^2 - 1)^\ell \right\}.$$

For all ℓ , μ_ℓ represents the ℓ -power sum of the mixed strategies of the players who mix and choose strategies from the set $\{1/k^2, \dots, 1/2\}$. Similarly, μ'_ℓ represents the ℓ -power sum of the mixed strategies of the players who mix and choose strategies from the set $\{1/2 + 1/k^2, \dots, (k^2 - 1)/k^2\}$. **Remark:** Whether there actually exist probability values $\pi_1, \dots, \pi_{t_s} \in \{1/k^2, \dots, 1/2\}$ and $\theta_1, \dots, \theta_{t_b} \in \{1/2 + 1/k^2, \dots, (k^2 - 1)/k^2\}$ such that $\mu_\ell = \sum_{i=1}^{t_s} \pi_i^\ell$ and $\mu'_\ell = \sum_{i=1}^{t_b} \theta_i^\ell$, for all $\ell = 1, 2, \dots, d$, will be determined later.

3. For each player $i = 1, \dots, n$, find a subset

$$\mathcal{S}_i \subseteq \left\{ 0, \frac{1}{k^2}, \dots, \frac{k^2 - 1}{k^2}, 1 \right\}$$

of permitted strategies for that player in an $\frac{\epsilon}{2}$ -Nash equilibrium, conditioned on the guesses in the previous steps. By this, we mean determining the answer to the following: “Given our

guesses for the aggregates $t_0, t_1, t_s, t_b, \mu_\ell, \mu'_\ell$, for all $\ell \in [d]$, what multiples of $1/k^2$ could player i be playing in an $\epsilon/2$ -Nash equilibrium?” Our test exploits the anonymity of the game and uses Theorem 7 to achieve the following:

- if a multiple of $1/k^2$ can be assigned to player i and complemented by choices of multiples for the other players, so that the aggregate conditions are satisfied and player i is at $3\epsilon/4$ -best response (that is, she experiences at most $3\epsilon/4$ regret), then this multiple of $1/k^2$ is included in the set \mathcal{S}_i ;
- if, given a multiple of $1/k^2$ to player i , there exists no assignment of multiples to the other players so that the aggregate conditions are satisfied and player i is at $3\epsilon/4$ -best response, the multiple is rejected from set \mathcal{S}_i .

Observe that the value of $3\epsilon/4$ for the regret used in the classifier is intentionally chosen midway between $\epsilon/2$ and ϵ . The reason for this value is that, if we only match the first d moments of a mixed strategy profile, our estimation of the real regret in that strategy profile is distorted by an additive error of $\epsilon/4$ (coming from (2) and the choice of d). Hence, with a threshold at $3\epsilon/4$ we make sure that: a. we are not going to “miss” the $\epsilon/2$ -Nash equilibrium (that we know exists in multiples of $1/k^2$ by virtue of our choice of a larger k), and b. *any* strategy profile that is consistent with the aggregate conditions and the sets \mathcal{S}_i found in this step is going to have regret at most $3\epsilon/4 + \epsilon/4 = \epsilon$. The fairly involved details of our test are given in Appendix D.2, and the way its analysis ties in with the search for an ϵ -Nash equilibrium is given in the proofs of Claims 9 and 10 of Appendix D.3.

4. Find an assignment of mixed strategies $v_1 \in \mathcal{S}_1, \dots, v_n \in \mathcal{S}_n$ to players, such that:

- t_0 players are assigned value 0 and t_1 players are assigned value 1;
- t_s players are assigned a value in $(0, 1/2]$ and $\sum_{i: v_i \in (0, 1/2]} v_i^\ell = \mu_\ell$, for all $\ell \in [d]$;
- t_b players are assigned a value in $(1/2, 1)$ and $\sum_{i: v_i \in (1/2, 1)} v_i^\ell = \mu'_\ell$, for all $\ell \in [d]$.

Solving this assignment problem is non-trivial, but it can be done by dynamic programming in time

$$O(n^3) \cdot \left(\frac{1}{\epsilon}\right)^{O(\log^2(1/\epsilon))},$$

because the sets \mathcal{S}_i are subsets of $\{0, 1/k^2, \dots, 1\}$. The algorithm is given in the proof of Claim 11 in Appendix D.3.

5. If an assignment is found, then the vector (v_1, \dots, v_n) constitutes an ϵ -Nash equilibrium.

Theorem 8. *MOMENT SEARCH is a PTAS for n -player 2-strategy anonymous games with running time $U \cdot \text{poly}(n) \cdot (1/\epsilon)^{O(\log^2(1/\epsilon))}$, where U is the number of bits required to represent a payoff value of the game. The algorithm generalizes to a constant number of player types with the number of types multiplying the exponent of the running time.*

Sketch: Correctness follows from this observation: The results in [7] and the choice of k guarantee that an $\frac{\epsilon}{2}$ -approximate Nash equilibrium in discretized probability values exists; therefore, Step 3 will find non-empty \mathcal{S}_i ’s for all players (for some guesses in Steps 1 and 2, since in particular the $\epsilon/2$ -Nash equilibrium will survive the tests of Step 3—by Theorem 7 and the choice of d , at most $\epsilon/4$ accuracy is lost if the correct values for the moments are guessed); and thus Step 4 will find an ϵ -approximate Nash equilibrium (another $\epsilon/4$ might be lost in this step). The full proof and the running time analysis are provided in Appendix D.3.

6 A Sparse ϵ -cover for Sums of Indicators

A byproduct of our proof is showing the existence of a sparse (and efficiently computable) ϵ -cover of the set of sums of independent indicators, under the total variation distance. To state our cover theorem, let $\mathcal{S} := \{\{X_i\}_i \mid X_1, \dots, X_n \text{ are independent indicators}\}$. We show the following.

Theorem 9 (Cover for sums of indicators). *For all $\epsilon > 0$, there exists a set $\mathcal{S}_\epsilon \subseteq \mathcal{S}$ such that (i) $|\mathcal{S}_\epsilon| \leq n^3 \cdot O(1/\epsilon) + n \cdot \left(\frac{1}{\epsilon}\right)^{O(\log^2 1/\epsilon)}$; (ii) For every $\{X_i\}_i \in \mathcal{S}$ there exists some $\{Y_i\}_i \in \mathcal{S}_\epsilon$ such that $d_{\text{TV}}(\sum_i X_i, \sum_i Y_i) \leq \epsilon$; and (iii) the set \mathcal{S}_ϵ can be constructed in time $O\left(n^3 \cdot O(1/\epsilon) + n \cdot \left(\frac{1}{\epsilon}\right)^{O(\log^2 1/\epsilon)}\right)$. Moreover, if $\{Y_i\}_i \in \mathcal{S}_\epsilon$, then the collection $\{Y_i\}_i$ has one of the following forms, where $k = k(\epsilon) = O(1/\epsilon)$ is a positive integer:*

- (Sparse Form) *There is a value $\ell \leq k^3 = O(1/\epsilon^3)$ such that for all $i \leq \ell$ we have $\mathcal{E}[Y_i] \in \left\{\frac{1}{k^2}, \frac{2}{k^2}, \dots, \frac{k^2-1}{k^2}\right\}$, and for all $i > \ell$ we have $\mathcal{E}[Y_i] \in \{0, 1\}$.*
- (k -heavy Binomial Form) *There is a value $\ell \in \{0, 1, \dots, n\}$ and a value $q \in \left\{\frac{1}{kn}, \frac{2}{kn}, \dots, \frac{kn-1}{kn}\right\}$ such that for all $i \leq \ell$ we have $\mathcal{E}[Y_i] = q$; for all $i > \ell$ we have $\mathcal{E}[Y_i] \in \{0, 1\}$; and ℓ, q satisfy the bounds $\ell q \geq k^2 - \frac{1}{k}$ and $\ell q(1 - q) \geq k^2 - k - 1 - \frac{3}{k}$.*

Proof of Theorem 9: Daskalakis [8] establishes the same theorem, except that the size of the cover he produces, as well as the time needed to produce it, are $n^3 \cdot O(1/\epsilon) + n \cdot \left(\frac{1}{\epsilon}\right)^{O(1/\epsilon^2)}$. Indeed, this bound is obtained by enumerating over all possible collections $\{Y_i\}_i$ in sparse form and all possible collections in k -heavy Binomial Form, for $k = O(1/\epsilon)$ specified by the theorem. Indeed, the total number of collections in k -heavy Binomial form is at most $(n+1)^2 nk = n^3 \cdot O(1/\epsilon)$, since there are at most $n+1$ choices for the value of ℓ , at most kn choices for the value of q , and at most $n+1$ choices for the number of variables indexed by $i > \ell$ that have expectation equal to 1 (the precise subset of these that have expectation 1 is not important, since this does not affect the distribution of $\sum_i Y_i$). On the other hand, the number of collections $\{Y_i\}_i$ in sparse form is at most $(k^3 + 1) \cdot k^{3k^2} \cdot (n+1) = n \cdot \left(\frac{1}{\epsilon}\right)^{O(1/\epsilon^2)}$, since there are $k^3 + 1$ choices for ℓ , k^{3k^2} choices for the expectations of variables Y_1, \dots, Y_ℓ up to permutations of the indices of these variables (namely we need to choose how many of these ℓ variables have expectation $1/k^2$, how many have expectation $2/k^2$, etc.), and at most $n+1$ choices for the number of variables indexed by $i > \ell$ that have expectation equal to 1.

To improve on the size of the cover we show that we can remove from the aforementioned cover a large fraction of collections in sparse form. In particular, we shall only keep $n \cdot \left(\frac{1}{\epsilon}\right)^{O(\log^2 1/\epsilon)}$ collections in sparse form, making use of Theorem 7. Indeed, consider a collection $\mathcal{Y} = \{Y_i\}_i$ in sparse form and let $\mathcal{L}_\mathcal{Y} = \{i \mid \mathcal{E}[Y_i] \in (0, 1/2]\} \subseteq [n]$, $\mathcal{R}_\mathcal{Y} = \{i \mid \mathcal{E}[Y_i] \in (1/2, 1)\} \subseteq [n]$. Theorem 7 implies that, if we compare $\{Y_i\}_i$ with another collection $\{Z_i\}_i$ satisfying the following:

$$\sum_{i \in \mathcal{L}} \mathcal{E}[Y_i]^t = \sum_{i \in \mathcal{L}} \mathcal{E}[Z_i]^t, \quad \text{for all } t = 1, \dots, d; \quad (3)$$

$$\sum_{i \in \mathcal{R}} \mathcal{E}[Y_i]^t = \sum_{i \in \mathcal{R}} \mathcal{E}[Z_i]^t, \quad \text{for all } t = 1, \dots, d; \quad (4)$$

$$\mathcal{E}[Y_i] = \mathcal{E}[Z_i], \quad \text{for all } i \in [n] \setminus (\mathcal{L} \cup \mathcal{R}), \quad (5)$$

then $d_{\text{TV}}(\sum_i Y_i, \sum_i Z_i) \leq 2 \cdot 20(d+1)^{1/4} 2^{-(d+1)/2}$. In particular, for some $d(\epsilon) = O(\log 1/\epsilon)$, this bound becomes at most ϵ .

For a collection $\mathcal{Y} = \{Y_i\}_i$, we define the *moment profile* $m_{\mathcal{Y}}$ of the collection to be a $(2d(\epsilon) + 1)$ -dimensional vector

$$m_{\mathcal{Y}} = \left(\sum_{i \in \mathcal{L}_{\mathcal{Y}}} \mathcal{E}[Y_i], \sum_{i \in \mathcal{L}_{\mathcal{Y}}} \mathcal{E}[Y_i]^2, \dots, \sum_{i \in \mathcal{L}_{\mathcal{Y}}} \mathcal{E}[Y_i]^{d(\epsilon)}; \sum_{i \in \mathcal{R}_{\mathcal{Y}}} \mathcal{E}[Y_i], \dots, \sum_{i \in \mathcal{R}_{\mathcal{Y}}} \mathcal{E}[Y_i]^{d(\epsilon)}; |\{i \mid \mathcal{E}[Y_i] = 1\}| \right).$$

By the previous discussion, for two collections $\mathcal{Y} = \{Y_i\}_i$ and $\mathcal{Z} = \{Z_i\}_i$, if $m_{\mathcal{Y}} = m_{\mathcal{Z}}$ then $d_{\text{TV}}(\sum_i Y_i, \sum_i Z_i) \leq \epsilon$.

Now given the ϵ -cover produced in [8] we perform the following sparsification operation: for every possible moment vector that can arise from a collection $\{Y_i\}_i$ in sparse form, we only keep in our cover one collection with such moment vector. The cover resulting from the sparsification operation is a 2ϵ -cover, since the sparsification loses us an additive ϵ in total variation distance, as argued above. We now compute the size of the new cover. The total number of moment vectors arising from sparse-form collections of indicators is at most $k^{O(d(\epsilon)^2)} \cdot (n + 1)$. Indeed, consider a collection \mathcal{Y} in sparse form. There are at most $k^3 + 1$ choices for $|\mathcal{L}_{\mathcal{Y}}|$, at most $k^3 + 1$ choices for $|\mathcal{R}_{\mathcal{Y}}|$, and at most $(n + 1)$ choices for $|\{i \mid \mathcal{E}[Y_i] = 1\}|$. We claim next that the total number of possible vectors

$$\left(\sum_{i \in \mathcal{L}_{\mathcal{Y}}} \mathcal{E}[Y_i], \sum_{i \in \mathcal{L}_{\mathcal{Y}}} \mathcal{E}[Y_i]^2, \dots, \sum_{i \in \mathcal{L}_{\mathcal{Y}}} \mathcal{E}[Y_i]^{d(\epsilon)} \right)$$

is at most $k^{O(d(\epsilon)^2)}$. Indeed, for all $t = 1, \dots, d(\epsilon)$, $\sum_{i \in \mathcal{L}_{\mathcal{Y}}} \mathcal{E}[Y_i]^t \leq |\mathcal{L}_{\mathcal{Y}}|$ and it must be a multiple of $1/k^{2t}$. So the total number of possible values for $\sum_{i \in \mathcal{L}_{\mathcal{Y}}} \mathcal{E}[Y_i]^t$ is at most $(k^{2t}|\mathcal{L}_{\mathcal{Y}}| + 1) \leq (k^{2t}k^3 + 1)$. It's easy to see then that the number of possible moment vectors

$$\left(\sum_{i \in \mathcal{L}_{\mathcal{Y}}} \mathcal{E}[Y_i], \sum_{i \in \mathcal{L}_{\mathcal{Y}}} \mathcal{E}[Y_i]^2, \dots, \sum_{i \in \mathcal{L}_{\mathcal{Y}}} \mathcal{E}[Y_i]^{d(\epsilon)} \right)$$

is at most

$$\prod_{t=1}^{d(\epsilon)} (k^{2t}k^3 + 1) \leq k^{O(d(\epsilon)^2)}.$$

The same upper bound applies to the total number of possible moment vectors

$$\left(\sum_{i \in \mathcal{R}_{\mathcal{Y}}} \mathcal{E}[Y_i], \sum_{i \in \mathcal{R}_{\mathcal{Y}}} \mathcal{E}[Y_i]^2, \dots, \sum_{i \in \mathcal{R}_{\mathcal{Y}}} \mathcal{E}[Y_i]^{d(\epsilon)} \right).$$

It follows then that the total number of sparse-form collections of indicators that we have kept in our cover after the sparsification operation is at most $k^{O(d(\epsilon)^2)} \cdot (n + 1) = n \cdot \left(\frac{1}{\epsilon}\right)^{O(\log^2 1/\epsilon)}$. The number of collections in heavy Binomial form that we have in our cover is the same as before and hence it is at most $n^3 \cdot O(1/\epsilon)$. So the size of the sparsified cover is at most $n^3 \cdot O(1/\epsilon) + n \cdot \left(\frac{1}{\epsilon}\right)^{O(\log^2 1/\epsilon)}$.

To finish the proof it remains to argue that we don't actually need to produce the cover of [8] and subsequently sparsify it to obtain our cover, but we can produce it directly in time $n^3 \cdot O(1/\epsilon) + n \cdot \left(\frac{1}{\epsilon}\right)^{O(\log^2 1/\epsilon)}$. We claim that given a moment vector m we can compute a collection $\mathcal{Y} = \{Y_i\}_i$ such that $m_{\mathcal{Y}} = m$, if such a collection exists, in time $\left(\frac{1}{\epsilon}\right)^{O(\log^2 1/\epsilon)}$. This follows from Claim 11 in

Appendix D.3.⁴ Hence, our algorithm producing the cover enumerates over all possible moment vectors and for each moment vector invokes Claim 11 to find a consistent sparse collection of indicators, if such collection exists, adding that collection into the cover. Then it enumerates over collections of indicators in heavy Binomial form and adds them to the cover. The overall running time is as promised. ■

7 Discussion and Open Problems

The mystery of PTAS for Nash equilibria deepens. There are simple algorithms for interesting special cases well within reach, and in fact we have seen that the existence of a PTAS is not incompatible with PPAD-completeness. But oblivious algorithms cannot take us all the way to the coveted PTAS for the general case. In the important special case of anonymous games, the approach of [13, 14, 7] — by design involving oblivious algorithms — hits a brick wall of $(\frac{1}{\epsilon})^{\frac{1}{\epsilon\alpha}}$, but then a more elaborate probabilistic result about moments and Bernoulli sums breaks that barrier. Pseudopolynomial bounds, familiar from [18], show up in approximation algorithms for anonymous games as well.

Many open problems remain, of course:

- Is there a PTAS for Nash equilibria in general games? A PTAS for bimatrix games that exploits the linear programming-like nature of the problem would not be unthinkable.
- Find a truly practical, and hopefully evocative of strategically interacting crowds, PTAS for anonymous games with two strategies.
- Prove that finding an exact Nash equilibrium in an anonymous game with a finite number of strategies is PPAD-complete.
- Find a PTAS for 2-strategy graphical games — the other important class of multi-player games.
- Alternatively, it is not unthinkable that the graphical games special case above is PPAD-complete to approximate sufficiently close.

References

- [1] I. Althöfer. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and Applications*, 199:339–355, 1994.
- [2] M. Blonski. Anonymous Games with Binary Actions. *Games and Economic Behavior*, 28(2): 171–180, 1999.
- [3] M. Blonski. The women of Cairo: Equilibria in large anonymous games. *Journal of Mathematical Economics*, 41(3): 253–264, 2005.
- [4] H. Bosse, J. Byrka and E. Markakis. New Algorithms for Approximate Nash Equilibria in Bimatrix Games. *WINE*, 2007.

⁴A naive application of Claim 11 results in running time $n^3 \cdot (\frac{1}{\epsilon})^{O(\log^2 1/\epsilon)}$. However, we can proceed as follows: we can guess $|\mathcal{L}_{\mathcal{Y}}|$ and $|\mathcal{R}_{\mathcal{Y}}|$ (at most $O(k^6)$ guesses) and invoke Claim 11 with $m_0 = m_1 = 0$, $m = 1$, $m_s = |\mathcal{L}_{\mathcal{Y}}|$ and $m_b = |\mathcal{R}_{\mathcal{Y}}|$ just to find $\{Y_i\}_{i \in \mathcal{L}_{\mathcal{Y}} \cup \mathcal{R}_{\mathcal{Y}}}$. To obtain the sought after $\mathcal{Y} = \{Y_i\}_i$ we then add $m_{2d(\epsilon)+1}$ indicators with expectation 1 and make the remaining indicators 0.

- [5] X. Chen and X. Deng. Settling the Complexity of Two-Player Nash Equilibrium. *FOCS*, 2006.
- [6] X. Chen, X. Deng and S.-H. Teng. Sparse Games Are Hard. *WINE*, 2006.
- [7] C. Daskalakis. An Efficient PTAS for Two-Strategy Anonymous Games. *WINE*, 2008.
- [8] C. Daskalakis. An Efficient PTAS for Two-Strategy Anonymous Games. *ArXiv Report*, 2008.
- [9] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The Complexity of Computing a Nash Equilibrium. *STOC*, 2006.
- [10] C. Daskalakis, A. Mehta, and C. H. Papadimitriou. A Note on Approximate Nash Equilibria. *WINE*, 2006.
- [11] C. Daskalakis, A. Mehta, and C. H. Papadimitriou. Progress in Approximate Nash Equilibria. *EC*, 2007.
- [12] C. Daskalakis and C. H. Papadimitriou. Three-Player Games Are Hard. *Electronic Colloquium in Computational Complexity*, TR05-139, 2005.
- [13] C. Daskalakis and C. H. Papadimitriou. Computing Equilibria in Anonymous Games. *FOCS*, 2007.
- [14] C. Daskalakis and C. H. Papadimitriou. Discretized Multinomial Distributions and Nash Equilibria in Anonymous Games. *FOCS*, 2008.
- [15] R. Durrett. *Probability: Theory and Examples*. Second Edition, Duxbury Press, 1996.
- [16] P. W. Goldberg and C. H. Papadimitriou. Reducibility Among Equilibrium Problems. *STOC*, 2006.
- [17] E. Kalai. Partially-Specified Large Games. *WINE*, 2005.
- [18] R. Lipton, E. Markakis, and A. Mehta. Playing Large Games Using Simple Strategies. *EC*, 2003.
- [19] I. Milchtaich. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior*, 13:111–124.
- [20] B. Roos. Binomial Approximation to the Poisson Binomial Distribution: The Krawtchouk Expansion. *Theory of Probability and its Applications*, 45(2):258–272, 2000.
- [21] H. Tsaknakis and P. G. Spirakis. An Optimization Approach for Approximate Nash Equilibria. *WINE*, 2007.
- [22] V. M. Zolotarev. Random Symmetric Polynomials. *Journal of Mathematical Sciences*, 38(5):2262–2272, 1987.

A Bimatrix Games

A.1 PTAS for Small Probability Games

Proof of Lemma 1: Let X and Y be independent random variables such that

$$X = e_i, \text{ with probability } x_i, \text{ for all } i \in [n],$$

$$Y = e_j, \text{ with probability } y_j, \text{ for all } j \in [n].$$

Let then X_1, X_2, \dots, X_t be t copies of variable X and Y_1, Y_2, \dots, Y_t be t copies of variable Y , where the variables $X_1, \dots, X_t, Y_1, \dots, Y_t$ are mutually independent. Setting $\mathcal{X} = \frac{1}{t} \sum_{k=1}^t X_k$ and $\mathcal{Y} = \frac{1}{t} \sum_{k=1}^t Y_k$, as in the statement of the theorem, we will argue that, with high probability, $(\mathcal{X}, \mathcal{Y})$ is an ϵ -Nash equilibrium of \mathcal{G} .

Let $\mathcal{U}_i = e_i^T R \mathcal{Y}$ and $\mathcal{V}_j = \mathcal{X}^T C e_j$ be the payoff of the row and column player respectively for playing strategy $i \in [n]$. Fixing $i \in [n]$, we have that

$$\mathcal{U}_i = \frac{1}{t} \sum_{k=1}^t e_i^T R Y_k,$$

$$\begin{aligned} \mathbb{E}[\mathcal{U}_i] &= \frac{1}{t} \sum_{k=1}^t \mathbb{E}[e_i^T R Y_k] = \frac{1}{t} \sum_{k=1}^t \mathbb{E}[e_i^T R Y] \\ &= \mathbb{E}[e_i^T R Y] = \sum_{j \in [n]} y_j e_i^T R e_j = e_i^T R y. \end{aligned}$$

An application of McDiarmid's inequality on the function $f(\Lambda_1, \dots, \Lambda_t) = \frac{1}{t} \sum_{k=1}^t \Lambda_k$, where $\Lambda_k := e_i^T R Y_k$, $k \in [t]$, are independent random variables, gives

$$\Pr[|\mathcal{U}_i - \mathbb{E}[\mathcal{U}_i]| \geq \epsilon/2] \leq 2e^{-\frac{t\epsilon^2}{8}} \leq \frac{2}{n^2}.$$

Applying a union bound, it follows that with probability at least $1 - \frac{4}{n}$ the following properties are satisfied by the pair of strategies $(\mathcal{X}, \mathcal{Y})$:

$$|e_i^T R \mathcal{Y} - e_i^T R y| \leq \epsilon/2, \text{ for all } i \in [n]; \quad (6)$$

$$|\mathcal{X}^T C e_j - x^T C e_j| \leq \epsilon/2, \text{ for all } j \in [n]. \quad (7)$$

The above imply that $(\mathcal{X}, \mathcal{Y})$ is an ϵ -Nash equilibrium. Indeed, for all $i, i' \in [n]$, we have that

$$\begin{aligned} e_i^T R \mathcal{Y} &> e_{i'}^T R \mathcal{Y} + \epsilon \\ &\Rightarrow e_i^T R y \geq e_i^T R \mathcal{Y} - \epsilon/2 > e_{i'}^T R \mathcal{Y} + \epsilon/2 > e_{i'}^T R y \\ &\Rightarrow x_{i'} = 0 \Rightarrow \mathcal{X}_{i'} = 0, \end{aligned}$$

where the last implication follows from the fact that \mathcal{X} is formed by taking a distribution over samples from x . Similarly, it can be argued that, for all $j, j' \in [n]$,

$$\mathcal{X}^T C e_j > \mathcal{X}^T C e_{j'} + \epsilon \Rightarrow \mathcal{Y}_{j'} = 0.$$

■

Proof of Claim 1: We are only going to lower bound the cardinality of the set \mathcal{A} . A similar argument applies for \mathcal{B} . Recall that (x, y) is a Nash equilibrium in which $x_i \leq \frac{1}{\delta n}$, for all $i \in [n]$, and similarly for y .

Let us now define the set \mathcal{A}' as follows

$$\mathcal{A}' := \{A \mid A \in \mathcal{A} \text{ and } x_i > 0, \forall i \in A\} \subseteq \mathcal{A},$$

that is, \mathcal{A}' is the subset of \mathcal{A} containing those multisets that could arise by taking t samples from x . We are going to use the structure of the probability distribution x and Lemma 1 to argue that \mathcal{A}' is large. For this, let us consider the random experiment of taking t independent samples from x and forming the corresponding multiset A . Let also \mathcal{X} be the uniform distribution over A . Lemma 1 asserts that with probability at least $1 - \frac{4}{n}$ the distribution \mathcal{X} will satisfy assertion 3 of Lemma 1. However, this does not directly imply a lower bound on the cardinality of \mathcal{A}' —it could be that all the probability mass is concentrated on a single element of \mathcal{A}' . However, the probability that a multiset arises by sampling x is at most

$$\left(\frac{1}{\delta n}\right)^t,$$

since the probability mass that x assigns to every $i \in [n]$ is at most $\frac{1}{\delta n}$. Therefore, the total number of distinct good multisets in \mathcal{A}' is at least

$$\frac{1 - \frac{4}{n}}{\left(\frac{1}{\delta n}\right)^t}.$$

Therefore, $|\mathcal{A}'| \geq (1 - \frac{4}{n})(\delta n)^t$ and $|\mathcal{A}| \geq |\mathcal{A}'| \geq (1 - \frac{4}{n})(\delta n)^t$. ■

Proof of Claim 2: From Claim 1, it follows that, if a random multiset A' is sampled, the probability that it is good is

$$\begin{aligned} \Pr[A' \in \mathcal{A}] &\geq \frac{(1 - \frac{4}{n})(\delta n)^t}{n^t} = \left(1 - \frac{4}{n}\right) \delta^t \\ &= \Omega\left(\delta \cdot n^{-16 \log(1/\delta)/\epsilon^2}\right), \end{aligned}$$

Similarly, if a random multiset B' is sampled, the probability that it is good is

$$\Pr[B' \in \mathcal{B}] = \Omega\left(\delta \cdot n^{-16 \log(1/\delta)/\epsilon^2}\right).$$

By independence, it follows that

$$\Pr[A' \in \mathcal{A} \text{ and } B' \in \mathcal{B}] = \Omega\left(\delta^2 \cdot n^{-32 \log(1/\delta)/\epsilon^2}\right).$$

But, if $A' \in \mathcal{A}$, $B' \in \mathcal{B}$, then the uniform distribution \mathcal{X}' over A' and the uniform distribution \mathcal{Y}' over B' comprise an ϵ -Nash equilibrium (see Proof of Lemma 1 for a justification). Hence,

$$\begin{aligned} \Pr[(\mathcal{X}', \mathcal{Y}') \text{ is an } \epsilon\text{-Nash equilibrium}] \\ = \Omega\left(\delta^2 \cdot n^{-32 \log(1/\delta)/\epsilon^2}\right). \end{aligned}$$

■

$$\begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 \\ \hline -1 & -1 & -1 & -1 & -1 & -1 \\ -1 & -1 & -1 & -1 & -1 & -1 \end{pmatrix}$$

Figure 1: R_S for the case $n = 6$, $\ell = 4$, $S = \{1, 2, 3, 4\}$.

A.2 The Lower Bound for Bimatrix Games

In Figure 1, we illustrate the game matrix R_S in our construction in the Proof of Theorem 4 (Section 3), for the case $n = 6$, $\ell = 4$, $S = \{1, 2, 3, 4\}$. Each column corresponds to a subset of S of size 2. We add then two extra rows to make the matrix square.

Proof of Lemma 2: The first assertion is easy to justify. Indeed, no matter what y is, the payoff u_i of the row player for playing strategy i , $i \in [n]$, satisfies:

- $u_i = -1$, for all $i \notin S$;
- $u_i \geq 0$, for all $i \in S$;

Hence, fixing any $i' \in S$, we have, for all $i \notin S$, that $u_{i'} > u_i + \epsilon$, which by the definition of an ϵ -Nash equilibrium implies that $x_i = 0$.

Towards justifying the second assertion, let u be the utility of the row player at the ϵ -Nash equilibrium (x, y) . Since, by the first assertion, $x_i = 0$ for all $i \notin S$, it follows that the payoff of the column player is $1 - u$, since the game restricted to the rows of the set S is 1-sum. Since every column of R_S restricted to the rows of the set S has exactly half of the entries equal to 1 and the other half equal to 0, it follows that

$$\sum_{i \in S} u_i = \sum_{j \in [n]} \frac{\ell}{2} y_j = \frac{\ell}{2}.$$

It follows that there exists some $i^* \in S$ such that $u_{i^*} \geq 1/2$. By the definition of ϵ -Nash equilibrium, it follows then that

$$u \geq u_{i^*} - \epsilon \geq 1/2 - \epsilon, \tag{8}$$

otherwise the row player would be including in his support strategies which are more than ϵ worse than i^* .

Let us now consider the payoff of the column player for choosing various strategies $j \in [n]$. Without loss of generality let us assume that $S = \{1, 2, \dots, \ell\}$ and that $x_1 \geq x_2 \geq \dots \geq x_\ell$. Let j^* be such that $S_{j^*} = \{\ell/2 + 1, \ell/2 + 2, \dots, \ell\}$. Then, by the definition of C_S , the payoff of the column player for choosing strategy j^* is

$$v_{j^*} = \sum_{i=1}^{\ell/2} x_i.$$

Since (x, y) is an ϵ -Nash equilibrium, it follows that v_{j^*} should be within ϵ from the payoff of the column player. Hence,

$$v_{j^*} \leq 1 - u + \epsilon.$$

Combining the above with (8), it follows that

$$v_{j^*} - \frac{1}{2} \leq 2\epsilon \Rightarrow \sum_{i=1}^{\ell/2} \left(x_i - \frac{1}{\ell} \right) \leq 2\epsilon.$$

The right hand side of the above, implies $\ell_1(x, u_S) \leq 8\epsilon$, by an application of Lemma 3 below, with $a_i = x_i - \frac{1}{\ell}$, for all $i \in [\ell]$, $k = 2\epsilon$.

Lemma 3. Let $\{a_i\}_{i=1}^{\ell}$ be real numbers satisfying the following properties for some $k \in \mathbb{R}_+$:

1. $a_1 \geq a_2 \geq \dots \geq a_{\ell}$;
2. $\sum_{i=1}^{\ell} a_i = 0$;
3. $\sum_{i=1}^{\ell/2} a_i \leq k$.

Then

$$\sum_{i=1}^{\ell} |a_i| \leq 4k. \quad (9)$$

Proof of Lemma 3: We distinguish two cases: (i) $a_{\ell/2} \geq 0$ and (ii) $a_{\ell/2} < 0$. In the first case, we have from Conditions 1 and 3 that

$$\sum_{i=1}^{\ell/2} |a_i| \leq k.$$

Using Condition 1 some more we get

$$\sum_{i > \frac{\ell}{2}: a_i \geq 0} |a_i| \leq \sum_{i > \frac{\ell}{2}: a_i \geq 0} a_{\ell/2} \leq \frac{\ell}{2} a_{\ell/2} \leq \sum_{i=1}^{\ell/2} |a_i| \leq k.$$

Combining the above we get

$$\sum_{i: a_i \geq 0} |a_i| \leq 2k. \quad (10)$$

Now we employ Condition 2, to get

$$\sum_{i: a_i < 0} |a_i| = \sum_{i: a_i < 0} (-a_i) = \sum_{i: a_i \geq 0} a_i = \sum_{i: a_i \geq 0} |a_i| \leq 2k. \quad (11)$$

We combine (10), (11) to deduce (9). Case (ii) is treated by repeating the above argument with

$$a_i \leftarrow (-a_{\ell-i+1}), \text{ for all } i \in [\ell].$$

■

■

Proof of Claim 3: Consider the set $V := \{u_S | S \subseteq [n], |S| = \ell\}$; note that $|V| = \Omega(n^{8 \log_2 n})$. For every $u_S \in V$ consider the set $N(u_S)$ of all other $u_{S'}$'s that are within ℓ_1 distance 17ϵ from u_S ; it is easy to see that $|N(u_S)| = O(n^{34\epsilon \log_2 n})$. Therefore, we can select a subset $V' \subseteq V$ of at least $\frac{|V|}{\max_S |N(u_S)|} = \Omega(n^{(.8-34\epsilon) \log_2 n})$ elements of V so that every pair of elements is at ℓ_1 distance at least 17ϵ apart; it follows that for every pair of elements in V' their ℓ_1 balls of radius 8ϵ are disjoint. The proof is complete. ■

B The Lower Bound for Oblivious PTAS's for Anonymous Games

B.1 Constructing Anonymous Games with Prescribed Equilibria

Theorem 10. *For any collection $\mathcal{P} := (p_i)_{i \in [k]}$, where $p_i \in [3\delta k, 1]$, for some $\delta > 0$ and $k \in \mathbb{N}$, there exists an anonymous game $\mathcal{G}_{\mathcal{P}}$ with $k+2$ players, 2 strategies, 0 and 1, payoffs in $[-1, 1]$, and three player types A, B, C , in which k players, $1, \dots, k$, belong to type A , 1 player belongs to type B , and 1 player belongs to type C , and such that in every δ' -Nash equilibrium, where $\delta' < \delta$, the following is satisfied: For every i , player i 's mixed strategy belongs to the set $[p_i - 7k^2\delta, p_i + 7k^2\delta]$; moreover, at least one of the players belonging to types B and C play strategy 1 with probability 0.*

Proof. Let us call B the single player of type B and C the single player of type C . Let us also use the notation: $\mu = \sum_{i \in [k]} p_i$, and $\mu_{-i} = \sum_{j \in [k] \setminus \{i\}} p_j$, for all i . Now, let us assign the following payoffs to the players B and C :

- $u_1^B = \frac{1}{k} \cdot (t_A - \mu)$, where t_A is the number of players of type A who play strategy 1;
- $u_0^B = 2\delta$;
- $u_1^C = \frac{1}{k} \cdot (\mu - t_A)$, where t_A is the number of players of type A who play strategy 1;
- $u_0^C = 2\delta$;

The payoff functions of the players of type A are defined as follows. For all $i \in [k]$:

- $u_0^i = \frac{1}{k}(\mu_{-i} \cdot \mathcal{X}_{B \text{ plays } 0} \cdot \mathcal{X}_{C \text{ plays } 0} - \delta k \cdot \mathcal{X}_{C \text{ plays } 1})$, where $\mathcal{X}_{B \text{ plays } 0}$, $\mathcal{X}_{C \text{ plays } 0}$ and $\mathcal{X}_{C \text{ plays } 1}$ are the indicators of the events ‘ B plays 0’, ‘ C plays 0’ and ‘ C plays 1’ respectively.
- $u_1^i = \frac{1}{k}(t_{A,-i} \cdot \mathcal{X}_{B \text{ plays } 0} \cdot \mathcal{X}_{C \text{ plays } 0} - \delta k \cdot \mathcal{X}_{B \text{ plays } 1})$, where $t_{A,-i}$ is the number of players of type A who are different than i and play 1, and $\mathcal{X}_{B \text{ plays } 0}$, $\mathcal{X}_{C \text{ plays } 0}$ and $\mathcal{X}_{B \text{ plays } 1}$ are the indicators of the events ‘ B plays 0’, ‘ C plays 0’ and ‘ B plays 1’ respectively.

Note that the range of all payoff functions of the game thus defined is $[-1, 1]$. We now claim the following:

Claim 4. *In every δ' -Nash equilibrium, where $\delta' < \delta$, it must be that*

$$\sum_{i \in [k]} q_i = \mu \pm 3\delta k,$$

where q_1, \dots, q_k are the mixed strategies of the players $1, \dots, k$.

Proof of Claim 4: Let $\mu' = \sum_{i \in [k]} q_i$. Suppose for a contradiction that in a δ' -Nash equilibrium $\mu' > \mu + 3\delta k$; then

$$\frac{1}{k}(\mu' - \mu) > 3\delta.$$

Note however that $\mathcal{E}[u_1^B] = \frac{1}{k}(\mu' - \mu)$ and $\mathcal{E}[u_1^C] = -\frac{1}{k}(\mu' - \mu)$. Hence, the above implies

$$\mathcal{E}[u_1^B] > \mathcal{E}[u_0^B] + \delta, \tag{12}$$

$$\mathcal{E}[u_1^C] < \mathcal{E}[u_0^C] - \delta. \tag{13}$$

Since $\delta' < \delta$, it must be then that $\Pr[B \text{ plays } 1] = 1$ and $\Pr[C \text{ plays } 1] = 0$. It follows then that for all $i \in [k]$:

$$\begin{aligned}\mathcal{E}[u_0^i] &= 0, \\ \mathcal{E}[u_1^i] &= -\delta.\end{aligned}$$

Hence, in a δ' -Nash equilibrium with $\delta' < \delta$, it must be that $\Pr[i \text{ plays } 1] = q_i = 0$, for all $i \in [k]$. This is a contradiction since we assumed that $\mu' = \sum_{i \in [k]} q_i > \mu + 3\delta k$, and μ is non-negative. Via similar arguments we show that the assumption $\mu' < \mu - 3\delta k$ also leads to a contradiction. Hence, in every δ' -Nash equilibrium with $\delta' < \delta$, it must be that

$$\mu' = \mu \pm 3\delta k.$$

■

We next show that in every δ' -Nash equilibrium with $\delta' < \delta$, at least one of the players B and C will not include strategy 1 in her support.

Claim 5. *In every δ' -Nash equilibrium with $\delta' < \delta$,*

$$\Pr[B \text{ plays } 1] = 0 \text{ or } \Pr[C \text{ plays } 1] = 0.$$

Proof of Claim 5: Let q_1, \dots, q_k be the mixed strategies of players $1, \dots, k$ at some δ' -Nash equilibrium of the game with $\delta' < \delta$. Let us consider the quantity $\mathcal{M} = \frac{1}{k}(\mu' - \mu)$, where $\mu' = \sum_{i \in [k]} q_i$. We distinguish the following cases:

- $\mathcal{M} \leq \delta$: In this case, $\mathcal{E}[u_1^B] = \frac{1}{k}(\mu' - \mu) \leq \delta \leq 2\delta - \delta = \mathcal{E}[u_0^B] - \delta$. Since $\delta' < \delta$, $\Pr[B \text{ plays } 1] = 0$.
- $\mathcal{M} \geq \delta$: In this case, $\mathcal{E}[u_1^C] = -\frac{1}{k}(\mu' - \mu) \leq -\delta \leq 2\delta - \delta = \mathcal{E}[u_0^C] - \delta$. Since $\delta' < \delta$, $\Pr[C \text{ plays } 1] = 0$.

■

Finally, we establish the following.

Claim 6. *In every δ' -Nash equilibrium with $\delta' < \delta$, it must be that for every player $i \in [k]$:*

$$\mu'_{-i} := \sum_{j \in [k] \setminus \{i\}} q_j = \mu_{-i} \pm 4\delta k^2,$$

where q_1, \dots, q_k are the mixed strategies of the players $1, \dots, k$.

Proof of Claim 6: Let us fix any δ' -Nash equilibrium. From Claim 5 it follows that either player B or C plays strategy 1 with probability 0. Without loss of generality, we will assume that $\Pr[C \text{ plays } 1] = 0$ (the argument for the case $\Pr[B \text{ plays } 1] = 0$ is identical to the one that follows).

Let us now fix a player $i \in [k]$. We show first that under the assumption $\Pr[C \text{ plays } 1] = 0$, $\Pr[C \text{ plays } 0] = 1$, it must be that

$$\mu_{-i} \leq \mu'_{-i} + \delta k. \tag{14}$$

Assume for a contradiction that $\mu_{-i} > \mu'_{-i} + \delta k$. It follows then that

$$\begin{aligned}
\mu_{-i}(1 - \Pr[\text{B plays 1}]) &\geq \mu'_{-i}(1 - \Pr[\text{B plays 1}]) \\
&\quad + \delta k(1 - \Pr[\text{B plays 1}]) \\
\Rightarrow \mu_{-i} \Pr[\text{B plays 0}] &\geq \mu'_{-i} \Pr[\text{B plays 0}] \\
&\quad + \delta k(1 - \Pr[\text{B plays 1}]) \\
\Rightarrow \mu_{-i} \Pr[\text{B plays 0}] \Pr[\text{C plays 0}] &\geq \\
&\quad \mu'_{-i} \Pr[\text{B plays 0}] \Pr[\text{C plays 0}] \\
&\quad - \delta k \Pr[\text{B plays 1}] + \delta k \\
\Rightarrow \mathcal{E}[u_0^i] &\geq \mathcal{E}[u_1^i] + \delta.
\end{aligned}$$

But we assumed that we fixed a δ' -Nash equilibrium with $\delta' < \delta$; hence the last equation implies that $q_i = 0$. But this leads quickly to a contradiction since, if $q_i = 0$, then using Claim 4 we have

$$\mu'_{-i} = \mu' \geq \mu - 3\delta k \geq \mu - p_i = \mu_{-i},$$

where we also used that $p_i \geq 3\delta k$. The above inequality contradicts our assumption that $\mu_{-i} > \mu'_{-i} + \delta k$. Hence, (14) must be satisfied. Using that $\mu' \leq \mu + 3\delta k$, which is implied by Claim 4) we get

$$q_i \leq p_i + 4\delta k.$$

From the above discussion it follows that

$$q_j \leq p_j + 4\delta k, \text{ for all } j. \quad (15)$$

Now fix $i \in [k]$ again. Summing (15) over all $j \neq i$, we get that

$$\mu'_{-i} \leq \mu_{-i} + 4\delta k^2. \quad (16)$$

Combining (14) and (16) we get

$$\mu'_{-i} = \mu_{-i} \pm 4\delta k^2.$$

■

To conclude the proof of Theorem 10, we combine Claims 4 and 6, as follows. For every player $i \in [k]$, we have from Claims 4 and 6 that in every δ' -Nash equilibrium with $\delta' < \delta$,

$$\mu'_{-i} = \mu_{-i} \pm 4\delta k^2 \text{ and } \mu' = \mu \pm 3\delta k.$$

By combining these equations we get

$$q_i = p_i \pm 7\delta k^2.$$

□

B.2 The Lower Bound

Given Theorem 10, we can establish our lower bound.

Proof of Theorem 5: Let us fix any oblivious ϵ -approximation algorithm for anonymous games with 2-strategies and 3-player types. The algorithm comes together with a distribution over unordered sets of mixed strategies—parametrized by the number of players n —which we denote by D_n .

We will consider the performance of the algorithm on the family of games specified in the statement of Theorem 10 for the following setting of parameters:

$$k = \lfloor (1/\epsilon)^{1/3} \rfloor, \delta = 1.01\epsilon, \mathcal{P} \in \mathcal{T}_\epsilon^k$$

$$\text{where } \mathcal{T}_\epsilon := \left\{ j \cdot 15\epsilon^{1/3} \mid j = 1, \dots, t_\epsilon \right\}, \quad t_\epsilon = \left\lfloor \frac{1}{15}\epsilon^{-1/3} \right\rfloor.$$

For technical reasons, let us define the following notion of distance between $\mathcal{P}, \mathcal{Q} \in \mathcal{T}_\epsilon^k$.

$$d(\mathcal{P}, \mathcal{Q}) := \sum_{j=1}^{t_\epsilon} |v_j^{\mathcal{P}} - v_j^{\mathcal{Q}}|.$$

where $v^{\mathcal{P}} = (v_1^{\mathcal{P}}, v_2^{\mathcal{P}}, \dots, v_{t_\epsilon}^{\mathcal{P}})$ is a vector storing the frequencies of various elements of the set \mathcal{T}_ϵ in the collection \mathcal{P} , i.e. $v_j^{\mathcal{P}} := |\{i \mid i \in [k], p_i = j \cdot 15\epsilon^{1/3}\}|$. To find the distance between two collections \mathcal{P}, \mathcal{Q} we compute the ℓ_1 distance of their frequency vectors. Notice in particular that this distance must be an even number. We also need the following definition.

Definition 2. We say that two anonymous games \mathcal{G} and \mathcal{G}' share an ϵ -Nash equilibrium in unordered form if there exists an ϵ -Nash equilibrium $\sigma_{\mathcal{G}}$ of game \mathcal{G} and an ϵ -Nash equilibrium $\sigma_{\mathcal{G}'}$ of game \mathcal{G}' such that $\sigma_{\mathcal{G}}$ and $\sigma_{\mathcal{G}'}$ are equal as unordered sets of mixed strategies.

We show first the following about the shareability of ϵ -Nash equilibria among the games $\mathcal{G}_{\mathcal{P}}, \mathcal{P} \in \mathcal{T}_\epsilon^k$.

Claim 7. If, for $\mathcal{P}, \mathcal{Q} \in \mathcal{T}_\epsilon^k$, $d(\mathcal{P}, \mathcal{Q}) > 0$, then there is no ϵ -Nash equilibrium that is shared between the games $\mathcal{G}_{\mathcal{P}}$ and $\mathcal{G}_{\mathcal{Q}}$ in unordered form.

Proof of Claim 7: For all j , let us define the $7.07k^2\epsilon$ ball around probability $j \cdot 15\epsilon^{1/3}$ in the natural way:

$$B_j := [j \cdot 15\epsilon^{1/3} - 7.07k^2\epsilon, j \cdot 15\epsilon^{1/3} + 7.07k^2\epsilon].$$

Observe that for all $j \geq 2$:

$$(j+1) \cdot 15\epsilon^{1/3} - j \cdot 15\epsilon^{1/3} = 15\epsilon^{1/3} > 2 \cdot 7.07k^2\epsilon.$$

Hence, for all j, j' : $B_j \cap B_{j'} = \emptyset$.

Now, let us consider any pair of ϵ -Nash equilibria $\sigma_{\mathcal{G}_{\mathcal{P}}}, \sigma_{\mathcal{G}_{\mathcal{Q}}}$ of the games $\mathcal{G}_{\mathcal{P}}$ and $\mathcal{G}_{\mathcal{Q}}$ and let us consider the vectors $v^{\sigma_{\mathcal{G}_{\mathcal{P}}}} = (v_1^{\sigma_{\mathcal{G}_{\mathcal{P}}}}, \dots, v_{t_\epsilon}^{\sigma_{\mathcal{G}_{\mathcal{P}}}})$ and $v^{\sigma_{\mathcal{G}_{\mathcal{Q}}}} = (v_1^{\sigma_{\mathcal{G}_{\mathcal{Q}}}}, \dots, v_{t_\epsilon}^{\sigma_{\mathcal{G}_{\mathcal{Q}}}})$ whose j -th components are defined as follows:

$$v_j^{\sigma_{\mathcal{G}_{\mathcal{P}}}} = \left(\begin{array}{c} \text{number of players who are} \\ \text{assigned a mixed strategy from} \\ \text{the set } B_j \text{ in } \sigma_{\mathcal{G}_{\mathcal{P}}} \end{array} \right),$$

$$v_j^{\sigma_{\mathcal{G}_{\mathcal{Q}}}} = \left(\begin{array}{c} \text{number of players who are} \\ \text{assigned a mixed strategy from} \\ \text{the set } B_j \text{ in } \sigma_{\mathcal{G}_{\mathcal{Q}}} \end{array} \right).$$

It is not hard to see that Theorem 10 and our assumption $d(\mathcal{P}, \mathcal{Q}) > 0$ imply that $\|v^{\sigma_{\mathcal{G}_{\mathcal{P}}}} - v^{\sigma_{\mathcal{G}_{\mathcal{Q}}}}\|_1 > 0$, hence $\sigma_{\mathcal{G}_{\mathcal{P}}}$ and $\sigma_{\mathcal{G}_{\mathcal{Q}}}$ cannot be permutations of each other. This concludes the proof. ■

Next, we show that there exists a large family of games such that no two members of the family share an ϵ -Nash equilibrium.

Claim 8. *There exists a subset $T \subseteq \mathcal{T}_\epsilon^k$ such that:*

1. *for every $\mathcal{P}, \mathcal{Q} \in T$: $d(\mathcal{P}, \mathcal{Q}) > 0$;*
2. $|T| \geq 2^{\Omega\left(\left(\frac{1}{\epsilon}\right)^{1/3}\right)}$;

Proof of claim 8: The total number of distinct multi-sets of cardinality k with elements from \mathcal{T}_ϵ is

$$\binom{t_\epsilon + k - 1}{k}.$$

Hence, it is easy to create a subset $T \subseteq \mathcal{T}_\epsilon^k$ such that:

- for every $\mathcal{P}, \mathcal{Q} \in T$: $d(\mathcal{P}, \mathcal{Q}) > 0$;
- $|T| = \binom{t_\epsilon + k - 1}{k}$.

Clearly, the set T satisfies Property 1 in the statement. For the cardinality bound we have:

$$\begin{aligned} |T| &\geq \binom{t_\epsilon + k - 1}{k} \geq \left(\frac{t_\epsilon + k - 1}{k}\right)^k \\ &\geq \left(1 + \frac{1}{15} - \frac{2}{k}\right)^k \geq 2^{\Omega\left(\left(\frac{1}{\epsilon}\right)^{1/3}\right)}. \end{aligned}$$

■

Now let us consider the performance of the distribution D_k on the family of anonymous games $\{\mathcal{G}_{\mathcal{P}}\}_{\mathcal{P} \in T}$, where T is the set defined in Claim 8. By Claims 7 and 8, no two games in the family share an ϵ -Nash equilibrium in unordered form. Hence, no matter what D_k is, there will be some game in our family for which the probability that D_k samples an ϵ -Nash equilibrium of that game is at most

$$1/|T| \leq 2^{-\Omega\left(\left(\frac{1}{\epsilon}\right)^{1/3}\right)}.$$

This concludes the proof of Theorem 5. ■

C The Binomial Approximation Theorem

Proposition 1. *Condition (C_d) in the statement of Theorem 7 is equivalent to the following condition:*

$$(V_d) : \quad \mathcal{E} \left[\left(\sum_{i=1}^n X_i \right)^\ell \right] = \mathcal{E} \left[\left(\sum_{i=1}^n Y_i \right)^\ell \right], \text{ for all } \ell \in [d].$$

Proof of Proposition 1:

$(V_d) \Rightarrow (C_d)$: It is not hard to see that $\mathcal{E} \left[\left(\sum_{i=1}^n X_i \right)^\ell \right]$ can be written as a weighted sum of the elementary symmetric polynomials $\psi_1(\mathcal{P}), \psi_2(\mathcal{P}), \dots, \psi_\ell(\mathcal{P})$ with positive coefficients, where, for all ℓ , $\psi_\ell(\mathcal{P})$ is defined as

$$\psi_\ell(\mathcal{P}) := \sum_{\substack{S \subseteq [n] \\ |S| = \ell}} \prod_{i \in S} p_i.$$

(V_d) then implies by induction

$$\psi_\ell(\mathcal{P}) = \psi_\ell(\mathcal{Q}), \quad \text{for all } \ell = 1, \dots, d. \quad (17)$$

Now, for all ℓ , define $\pi_\ell(\mathcal{P})$ as the power sum symmetric polynomial of degree ℓ

$$\pi_\ell(\mathcal{P}) := \sum_{i=1}^n p_i^\ell.$$

Now fix any $\ell \leq d$. Observe that $\pi_\ell(\mathcal{P})$ is a symmetric polynomial of degree ℓ on the variables p_1, \dots, p_n . It follows (see, e.g., [22]) that $\pi_\ell(\mathcal{P})$ can be expressed as a function of the elementary symmetric polynomials $\psi_1(\mathcal{P}), \dots, \psi_\ell(\mathcal{P})$. Since, by (17), $\psi_j(\mathcal{P}) = \psi_j(\mathcal{Q})$, for all $j \leq \ell$, we get that $\pi_\ell(\mathcal{P}) = \pi_\ell(\mathcal{Q})$. Since this holds for any $\ell \leq d$, (C_d) is satisfied.

The implication $(C_d) \Rightarrow (V_d)$ is established in a similar fashion. (C_d) implies

$$\pi_\ell(\mathcal{P}) = \pi_\ell(\mathcal{Q}), \text{ for all } \ell = 1, \dots, d.$$

For any $\ell \leq d$, $\mathcal{E} \left[\left(\sum_{i=1}^n X_i \right)^\ell \right]$ is a symmetric polynomial of degree ℓ on the variables p_1, \dots, p_n . It follows (see, e.g., [22]) that $\mathcal{E} \left[\left(\sum_{i=1}^n X_i \right)^\ell \right]$ can be expressed as a function of the power-sum symmetric polynomials $\pi_1(\mathcal{P}), \dots, \pi_\ell(\mathcal{P})$. And since $\pi_j(\mathcal{P}) = \pi_j(\mathcal{Q})$, for all $j \leq \ell$, we get $\mathcal{E} \left[\left(\sum_{i=1}^n X_i \right)^\ell \right] = \mathcal{E} \left[\left(\sum_{i=1}^n Y_i \right)^\ell \right]$. Since this holds for any $\ell \leq d$, (V_d) is satisfied. ■

Proof of Theorem 7: Let $X = \sum_i X_i$. The following theorem due to Roos [20], specifies an expansion of the distribution function of X as a sum of a finite number of signed measures: the binomial distribution $\mathcal{B}_{n,p}(m)$ (for an arbitrary choice of p) and its first n derivatives with respect to the parameter p , at the chosen value of p . More precisely,

Theorem 11 ([20]). *Let $\mathcal{P} := (p_i)_{i=1}^n$ be an arbitrary set of probability values in $[0, 1]$ and $\mathcal{X} := (X_i)_{i=1}^n$ a collection of independent indicators with $\mathcal{E}[X_i] = p_i$, for all $i \in [n]$; also let $X := \sum_i X_i$. Then, for all $m \in \{0, \dots, n\}$ and any $p \in (0, 1)$,*

$$\Pr[X = m] = \sum_{\ell=0}^n \alpha_\ell(\mathcal{P}, p) \cdot \delta^\ell \mathcal{B}_{n,p}(m), \quad (18)$$

where in the above $\alpha_0(\mathcal{P}, p) := 1$,

$$\alpha_\ell(\mathcal{P}, p) := \sum_{1 \leq k(1) < \dots < k(\ell) \leq n} \prod_{r=1}^{\ell} (p_{k(r)} - p), \quad \text{for all } \ell \in [n], \quad (19)$$

and

$$\delta^\ell \mathcal{B}_{n,p}(m) := \frac{(n-\ell)!}{n!} \frac{\partial^\ell}{\partial p^\ell} \mathcal{B}_{n,p}(m),$$

and for the purposes of the last definition we interpret

$$\mathcal{B}_{n,p}(m) := b(m, n, p)$$

as a function of the arguments m, n, p in the natural way:

$$b(m, n, p) := \begin{cases} \binom{n}{m} p^m (1-p)^{n-m}, & \text{for } n, m \in \mathbb{Z}_+, \ m \leq n; \\ 0, & \text{otherwise.} \end{cases}$$

Assume now that we are given two collections \mathcal{X}, \mathcal{Y} of indicators as in the statement of Theorem 7. We claim the following

Lemma 4. *For sets \mathcal{P} and \mathcal{Q} satisfying property (C_d) , and for all p :*

$$\alpha_\ell(\mathcal{P}, p) = \alpha_\ell(\mathcal{Q}, p), \quad \text{for all } \ell = 0, \dots, d.$$

Proof of lemma 4: Clearly, $\alpha_0(\mathcal{P}, p) = \alpha_0(\mathcal{Q}, p)$. Let us fix $\ell \in \{1, \dots, d\}$. Observe that $\alpha_\ell(\mathcal{P}, p)$ is a symmetric polynomial of degree ℓ on the variables p_1, \dots, p_n . It follows (see, e.g., [22]) that $\alpha_\ell(\mathcal{P}, p)$ can be expressed as a function of the power-sum symmetric polynomials $\pi_1(\mathcal{P}), \dots, \pi_\ell(\mathcal{P})$ defined as

$$\pi_j(\mathcal{P}) := \sum_{i=1}^n p_i^j, \quad \text{for all } j \in [\ell].$$

It follows from (C_d) that $\pi_j(\mathcal{P}) = \pi_j(\mathcal{Q})$, for all $j \leq \ell$; from the previous discussion, this implies that $\alpha_\ell(\mathcal{P}, p) = \alpha_\ell(\mathcal{Q}, p)$. ■

For any $p \in (0, 1)$, by combining Lemma 4 and Theorem 11 we get that

$$\begin{aligned} & Pr[X = m] - Pr[Y = m] \\ &= \sum_{\ell=d+1}^n (\alpha_\ell(\mathcal{P}, p) - \alpha_\ell(\mathcal{Q}, p)) \cdot \delta^\ell \mathcal{B}_{n,p}(m), \quad \text{for all } m. \end{aligned}$$

Therefore, for all p :

$$\|X; Y\| \leq \frac{1}{2} \sum_{m=0}^n |Pr[X = m] - Pr[Y = m]| \tag{20}$$

$$\leq \frac{1}{2} \sum_{\ell=d+1}^n |\alpha_\ell(\mathcal{P}, p) - \alpha_\ell(\mathcal{Q}, p)| \cdot \|\delta^\ell \mathcal{B}_{n,p}(\cdot)\|_1 \tag{21}$$

$$\leq \frac{1}{2} \sum_{\ell=d+1}^n (|\alpha_\ell(\mathcal{P}, p)| + |\alpha_\ell(\mathcal{Q}, p)|) \cdot \|\delta^\ell \mathcal{B}_{n,p}(\cdot)\|_1. \tag{22}$$

From Theorem 2 in [20], it follows that

$$\begin{aligned} & \frac{1}{2} \sum_{\ell=d+1}^n |\alpha_\ell(\mathcal{P}, p)| \cdot \|\delta^\ell \mathcal{B}_{n,p}(\cdot)\|_1 \\ & \leq \frac{\sqrt{e}(d+1)^{1/4}}{2} \theta(\mathcal{P}, p)^{(d+1)/2} \frac{1 - \frac{d}{d+1} \sqrt{\theta(\mathcal{P}, p)}}{(1 - \sqrt{\theta(\mathcal{P}, p)})^2}, \end{aligned}$$

where

$$\theta(\mathcal{P}, p) := \frac{2 \sum_{i=1}^n (p_i - p)^2 + (\sum_{i=1}^n (p_i - p))^2}{2np(1-p)}$$

Choosing $p = \bar{p} := \frac{1}{n} \sum_i p_i$, we get (see [20])

$$\theta(\mathcal{P}, \bar{p}) = \frac{\sum_{i=1}^n (p_i - \bar{p})^2}{n\bar{p}(1-\bar{p})} \leq |p_{\max} - p_{\min}| \leq \frac{1}{2},$$

where $p_{\max} = \max_i \{p_i\}$ and $p_{\min} = \min_i \{p_i\}$. From the above we get

$$\begin{aligned} \frac{1}{2} \sum_{\ell=d+1}^n |\alpha_\ell(\mathcal{P}, \bar{p})| \cdot \|\delta^\ell \mathcal{B}_{n, \bar{p}}(\cdot)\|_1 \\ \leq \sqrt{e}(d+1)^{1/4} 2^{-(d+1)/2} \frac{1 - \frac{1}{\sqrt{2}} \frac{d}{d+1}}{(\sqrt{2} - 1)^2} \\ \leq 10(d+1)^{1/4} 2^{-(d+1)/2}. \end{aligned}$$

Since from (C_d) we have that $\bar{p} = \sum_i p_i = \sum_i q_i$, we get in a similar fashion

$$\frac{1}{2} \sum_{\ell=d+1}^n |\alpha_\ell(\mathcal{Q}, \bar{p})| \cdot \|\delta^\ell \mathcal{B}_{n, \bar{p}}(\cdot)\|_1 \leq 10(d+1)^{1/4} 2^{-(d+1)/2}.$$

Plugging these bounds into (22) we get

$$\|X; Y\| \leq 20(d+1)^{1/4} 2^{-(d+1)/2}.$$

■

Corollary 1. *Let $\mathcal{P} := (p_i)_{i=1}^n \in [1/2, 1]^n$ and $\mathcal{Q} := (q_i)_{i=1}^n \in [1/2, 1]^n$ be two collections of probability values in $[1/2, 1]$. Let also $\mathcal{X} := (X_i)_{i=1}^n$ and $\mathcal{Y} := (Y_i)_{i=1}^n$ be two collections of independent indicators with $\mathcal{E}[X_i] = p_i$ and $\mathcal{E}[Y_i] = q_i$, for all $i \in [n]$. If for some $d \in [n]$ Condition (C_d) is satisfied, then*

$$\left\| \sum_i X_i ; \sum_i Y_i \right\| = 20(d+1)^{1/4} 2^{-(d+1)/2}.$$

Proof of Corollary 1: Define $X'_i = 1 - X_i$, and $Y'_i = 1 - Y_i$, for all i . Apply Theorem 7 to deduce

$$\left\| \sum_i X'_i ; \sum_i Y'_i \right\| = 20(d+1)^{1/4} 2^{-(d+1)/2}.$$

The proof is completed by noting

$$\left\| \sum_i X_i ; \sum_i Y_i \right\| = \left\| \sum_i X'_i ; \sum_i Y'_i \right\|.$$

■

D The non-oblivious PTAS for Anonymous Games

D.1 The Oblivious PTAS of [Daskalakis, 2008]

In [7], Theorem 6 was used to design an oblivious PTAS for n -player 2-strategy anonymous games running in time

$$\text{poly}(n) \cdot (1/\epsilon)^{O(1/\epsilon^2)} \cdot U,$$

where U is the number of bits required to represent a payoff value of the game. The algorithm has the following structure

1. choose $k = O(1/\epsilon)$, according to Theorem 6;
2. guess the number t of players who randomize, the number t_0 of players who play 0, and the number $t_1 = n - t - t_0$ of players who play 1;
3. depending on the number t of players who mix try one of the following:
 - (a) if $t \leq k^3$, guess the number of players ψ_i playing each of the integer multiples i/k^2 of $\frac{1}{k^2}$, and, solving a max-flow instance (see details in [7]), check if there is a ϵ -Nash equilibrium in which ψ_i players play i/k^2 , for all i , t_0 players play 0, and t_1 players play 1;
 - (b) if $t > k^3$, guess an integer multiple i/kn of $1/kn$ and, solving a max-flow instance, check if there is an ϵ -Nash equilibrium in which t players play i/kn , t_0 players play 0, and t_1 players play 1.

Figure 2: The oblivious PTAS of [Daskalakis, 2008]

Clearly, there are $O(n^2)$ possible choices for Step 2 of the algorithm. Moreover, the search of Step 3b can be completed in time (see [7])

$$U \cdot \text{poly}(n) \cdot (1/\epsilon) \log_2(1/\epsilon),$$

which is polynomial in $1/\epsilon$. On the other hand, Step 3a involves searching over all partitions of t balls into $k^2 - 2$ bins. The resulting running time for this step (see details in [7]) is

$$U \cdot \text{poly}(n) \cdot (1/\epsilon)^{O(1/\epsilon^2)},$$

which is exponential in $1/\epsilon$.

D.2 MOMENT SEARCH : Missing Details

We describe in detail the third step of MOMENT SEARCH.

3. For each player $i = 1, \dots, n$, find a subset

$$\mathcal{S}_i \subseteq \left\{ 0, \frac{1}{k^2}, \dots, \frac{k^2 - 1}{k^2}, 1 \right\}$$

of permitted mixed strategies for that player in an ϵ -Nash equilibrium, “conditioning” on the total number of players playing 0 being t_0 , the total number of players playing 1 being t_1 , and the probabilities of the players who mix resulting in the power-sums μ_1, \dots, μ_d and μ'_1, \dots, μ'_d . The way we compute the set \mathcal{S}_i is as follows:

- (a) To determine whether $0 \in \mathcal{S}_i$:
 - i. Find *any* set of mixed strategies $q_1, \dots, q_{t_s} \subseteq \{\frac{1}{k^2}, \frac{2}{k^2}, \dots, \frac{1}{2}\}$ such that $\sum_{i=1}^{t_s} q_i^\ell = \mu_\ell$, for all $\ell = 1, \dots, d$. Find *any* set of mixed strategies $r_1, \dots, r_{t_b} \subseteq \{\frac{1}{2} + \frac{1}{k^2}, \frac{1}{2} + \frac{2}{k^2}, \dots, 1 - \frac{1}{k^2}\}$ such that $\sum_{i=1}^{t_b} r_i^\ell = \mu'_\ell$, for all $\ell = 1, \dots, d$. If such values do not exist FAIL.

Remark: An efficient algorithm to solve this optimization problem is described in the proof of Claim 11.

ii. Define the random variable

$$Y = (t_0 - 1) \cdot 0 + \sum_{i=1}^{t_s} S_i + \sum_{i=1}^{t_b} B_i + t_1 \cdot 1,$$

where the variables $S_1, \dots, S_{t_s}, B_1, \dots, B_{t_b}$ are mutually independent with expectations $\mathcal{E}[S_i] = q_i$, for all $i = 1, \dots, t_s$, and $\mathcal{E}[B_i] = r_i$, for all $i = 1, \dots, t_b$.

iii. Compute the expected payoff $\mathcal{U}_0^i = \mathcal{E}[u_0^i(Y)]$ and $\mathcal{U}_1^i = \mathcal{E}[u_1^i(Y)]$ of player i for playing 0 and 1 respectively, if the number of the other players playing 1 is distributed identically to Y .

iv. if $\mathcal{U}_0^i \geq \mathcal{U}_1^i - 3\epsilon/4$, then include 0 to the set \mathcal{S}_i , otherwise do not.

(b) To determine whether $1 \in \mathcal{S}_i$, follow the same strategy except now Y is defined as follows

$$Y = t_0 \cdot 0 + \sum_{i=1}^{t_s} S_i + \sum_{i=1}^{t_b} B_i + (t_1 - 1) \cdot 1,$$

to account for the fact that we are testing for the candidate strategy 1 for player i . Also, the test that determines whether $1 \in \mathcal{S}_i$ is now $\mathcal{U}_1^i \geq \mathcal{U}_0^i - 3\epsilon/4$.

(c) For all $j \in \{1, \dots, k^2/2\}$, to determine whether $j/k^2 \in \mathcal{S}_i$ do the following slightly updated test:

i. Find any set of mixed strategies $q_1, \dots, q_{t_s-1} \subseteq \{\frac{1}{k^2}, \frac{2}{k^2}, \dots, 1/2\}$ such that $\sum_{i=1}^{t_s-1} q_i^\ell = \mu_\ell - (j/k^2)^\ell$, for all $\ell = 1, \dots, d$. Find any set of mixed strategies

$$r_1, \dots, r_{t_b} \subseteq \{\frac{1}{2} + \frac{1}{k^2}, \frac{1}{2} + \frac{2}{k^2}, \dots, 1 - \frac{1}{k^2}\}$$

such that $\sum_{i=1}^{t_b} r_i^\ell = \mu'_\ell$, for all $\ell = 1, \dots, d$. If such values do not exist FAIL.

ii. Define the random variable

$$Y = t_0 \cdot 0 + \sum_{i=1}^{t_s-1} S_i + \sum_{i=1}^{t_b} B_i + t_1 \cdot 1,$$

where the variables $S_1, \dots, S_{t_s-1}, B_1, \dots, B_{t_b}$ are mutually independent with $\mathcal{E}[S_i] = q_i$, for all $i = 1, \dots, t_s - 1$, and $\mathcal{E}[B_i] = r_i$, for all $i = 1, \dots, t_b$.

iii. Compute the expected payoff $\mathcal{U}_0^i = \mathcal{E}[u_0^i(Y)]$ and $\mathcal{U}_1^i = \mathcal{E}[u_1^i(Y)]$ of player i for playing 0 and 1 respectively, if the number of the other players playing 1 is distributed identically to Y .

iv. if $\mathcal{U}_0^i \in [\mathcal{U}_1^i - 3\epsilon/4, \mathcal{U}_1^i + 3\epsilon/4]$, then include j/k^2 to the set \mathcal{S}_i , otherwise do not.

(d) For all $j \in \{(k^2 + 2)/2, \dots, k^2 - 1\}$, to determine whether $j/k^2 \in \mathcal{S}_i$ do the appropriate modifications to the method described in Step 3c.

D.3 The Analysis of MOMENT SEARCH

Correctness The correctness of MOMENT SEARCH follows from the following two claims.

Claim 9. *If there exists an $\epsilon/2$ -Nash equilibrium in which $t \leq k^3$ players mix, and their mixed strategies are integer multiples of $1/k^2$, then MOMENT SEARCH will not fail, i.e. it will output a set of mixed strategies (v_1, \dots, v_n) .*

Claim 10. *If MOMENT SEARCH outputs a set of mixed strategies (v_1, \dots, v_n) , then these strategies constitute an ϵ -Nash equilibrium.*

Proof of Claim 9: Let (p_1, \dots, p_n) be an $\epsilon/2$ -Nash equilibrium in which t_0 players play 0, t_1 players play 1, and $t \leq k^3$ players mix, and their mixed strategies are integer multiples of $1/k^2$. It suffices to show that there exist guesses for $t_0, t_1, t_s, t_b, \mu_1, \dots, \mu_d, \mu'_1, \dots, \mu'_d$, such that $p_1 \in \mathcal{S}_1, p_2 \in \mathcal{S}_2, \dots, p_n \in \mathcal{S}_n$. Indeed, let

$$\mathcal{I}_0 := \{i | p_i = 0\}, \quad \mathcal{I}_s := \{i | p_i \in (0, 1/2]\},$$

$$\mathcal{I}_b := \{i | p_i \in (1/2, 1)\}, \quad \mathcal{I}_1 := \{i | p_i = 1\},$$

and let us choose the following values for our guesses

$$t_0 := |\mathcal{I}_0|, t_s = |\mathcal{I}_s|, t_b = |\mathcal{I}_b|, t_1 := |\mathcal{I}_1|$$

and, for all $\ell \in [d]$,

$$\mu_\ell = \sum_{i \in \mathcal{I}_s} p_i^\ell, \quad \mu'_\ell = \sum_{i \in \mathcal{I}_b} p_i^\ell.$$

We will show that for the guesses that we defined above $p_i \in \mathcal{S}_i$, for all i . We distinguish the following cases: $i \in \mathcal{I}_0, i \in \mathcal{I}_s, i \in \mathcal{I}_b, i \in \mathcal{I}_1$. The proof for all the cases proceeds in the same fashion. We will only argue about the case $i \in \mathcal{I}_s$; in particular, we will show that in Step 3c of MOMENT SEARCH the test succeeds for $j/k^2 = p_i$.

At the equilibrium point (p_1, \dots, p_n) , the number of the other players who choose strategy 1, from the perspective of player i , is distributed identically to the random variable:

$$Z := \sum_{j \in \mathcal{I}_s \setminus \{i\}} X_j + \sum_{j \in \mathcal{I}_b} X_j + t_1 \cdot 1,$$

where $\mathcal{E}[X_j] = p_j$ for all j . Since (p_1, \dots, p_n) is an $\epsilon/2$ -Nash equilibrium it must be the case that

$$|\mathcal{E}[u_0^i(Z)] - \mathcal{E}[u_1^i(Z)]| \leq \epsilon/2. \quad (23)$$

We will argue that, if in the above equation, we replace Z by Y , where Y is the random variable defined in Step 3(c)ii of MOMENT SEARCH, the inequality still holds with slightly updated upper bound:

$$|\mathcal{E}[u_0^i(Y)] - \mathcal{E}[u_1^i(Y)]| \leq 3\epsilon/4. \quad (24)$$

If (24) is established, the proof is completed since Step 3(c)iv will include j/k^2 into the set \mathcal{S}_i .

Let $S_1, \dots, S_{t_s-1}, B_1, \dots, B_{t_b}$ be the random variables with expectations $q_1, \dots, q_{t_s-1}, r_1, \dots, r_{t_b}$ defined in Step 3(c)ii of MOMENT SEARCH. Observe that, for all $\ell = 1, \dots, d$,

$$\sum_{j=1}^{t_s-1} q_j^\ell = \mu_\ell - (j/k^2)^\ell = \sum_{j \in \mathcal{I}_s \setminus \{i\}} p_j^\ell,$$

since $p_i = j/k^2$. Hence, by Theorem 7,

$$\left\| \sum_{j=1}^{t_s-1} S_j - \sum_{j \in \mathcal{I}_s \setminus \{i\}} X_j \right\| \leq 20(d+1)^{1/4} 2^{-(d+1)/2} \leq \epsilon/16. \quad (25)$$

Via similar arguments and Corollary 1, we get

$$\left\| \sum_{j=1}^{t_b} B_j - \sum_{j \in \mathcal{I}_b} X_j \right\| \leq \epsilon/16. \quad (26)$$

(25) and (26) imply via the coupling lemma

$$\|Y; Z\| \leq \frac{\epsilon}{8}. \quad (27)$$

It is not hard to see that

$$|\mathcal{E}[u_0^i(Y)] - \mathcal{E}[u_0^i(Z)]| \leq \|Y; Z\| \leq \frac{\epsilon}{8},$$

where we used (27). Similarly,

$$|\mathcal{E}[u_1^i(Y)] - \mathcal{E}[u_1^i(Z)]| \leq \|Y; Z\| \leq \frac{\epsilon}{8}.$$

Combining the above with (23) we get (24). This concludes the proof. ■

Proof of Claim 10: Let

$$\mathcal{I}_0 := \{i | v_i = 0\}, \quad \mathcal{I}_s := \{i | v_i \in (0, 1/2]\},$$

$$\mathcal{I}_b := \{i | v_i \in (1/2, 1)\}, \quad \mathcal{I}_1 := \{i | v_i = 1\},$$

$$t_s = |\mathcal{I}_s|, \quad t_b = |\mathcal{I}_b|,$$

Observe that the moment values that were guessed in Step 2 of MOMENT SEARCH satisfy

$$\mu_\ell = \sum_{i \in \mathcal{I}_s} v_i^\ell, \quad \mu'_\ell = \sum_{i \in \mathcal{I}_b} v_i^\ell, \quad \text{for all } \ell = 1, \dots, d.$$

We will argue that (v_1, \dots, v_n) is an ϵ -Nash equilibrium. To do this we need to argue that, for each player i , v_i is an ϵ -well supported strategy against the strategies of her opponents. We distinguish the following cases: $i \in \mathcal{I}_0$, $i \in \mathcal{I}_s$, $i \in \mathcal{I}_b$, $i \in \mathcal{I}_1$. The proof for all the cases proceeds in a similar fashion. We will only present the argument for the case $i \in \mathcal{I}_s$.

Let $v_i = j/k^2$ for some $j \in \{1, \dots, \frac{k^2}{2}\}$. From the perspective of player i , the number of other players who play 1 in the mixed strategy profile (v_1, \dots, v_n) is distributed identically to the random variable

$$Z := \sum_{j \in [n] \setminus \{i\}} X_j,$$

where $\mathcal{E}[X_j] = v_j$ for all j . To argue that v_i is an ϵ -well supported strategy against the strategies of i 's opponents, we need to show that

$$|\mathcal{E}[u_0^i(Z)] - \mathcal{E}[u_1^i(Z)]| \leq \epsilon. \quad (28)$$

Let us now go back to the iteration of Step 3c in which the probability value j/k^2 was inserted into the set \mathcal{S}_i . Let q_1, \dots, q_{t_s-1} , r_1, \dots, r_{t_b} be the values that were selected at Step 3(c)i of that iteration, and let

$$Y = \sum_{j=1}^{t_s-1} S_j + \sum_{j=1}^{t_b} B_j + t_1 \cdot 1,$$

be the random variable defined in Step 3(c)ii, where the variables $S_1, \dots, S_{t_s-1}, B_1, \dots, B_{t_b}$ are mutually independent with expectations $\mathcal{E}[S_i] = q_i$, for all $i = 1, \dots, t_s - 1$, and $\mathcal{E}[B_i] = r_i$, for all $i = 1, \dots, t_b$. Observe that the q_j 's and r_j 's were chosen by Step 3(c)i so that the following are satisfied

$$\sum_{j=1}^{t_s-1} q_j^\ell = \mu_\ell - (j/k^2)^\ell = \mu_\ell - v_i^\ell = \sum_{j \in \mathcal{I}_s \setminus \{i\}} v_j^\ell, \text{ for all } \ell \in [d], \quad (29)$$

$$\text{and } \sum_{j=1}^{t_b} r_j^\ell = \mu'_\ell = \sum_{j \in \mathcal{I}_b} v_j^\ell, \text{ for all } \ell = 1, \dots, d. \quad (30)$$

Equation (29) implies via Theorem 7 that

$$\left\| \sum_{j=1}^{t_s-1} S_j - \sum_{j \in \mathcal{I}_s \setminus \{i\}} X_j \right\| \leq 20(d+1)^{1/4} 2^{-(d+1)/2} \leq \epsilon/16. \quad (31)$$

Equation (30) and Corollary 1 imply

$$\left\| \sum_{j=1}^{t_b} B_j - \sum_{j \in \mathcal{I}_b} X_j \right\| \leq \epsilon/16. \quad (32)$$

(31) and (32) imply via the coupling lemma

$$\|Y; Z\| \leq \frac{\epsilon}{8}. \quad (33)$$

It is not hard to see that

$$|\mathcal{E}[u_0^i(Y)] - \mathcal{E}[u_0^i(Z)]| \leq \|Y; Z\| \leq \frac{\epsilon}{8}, \quad (34)$$

where we used (33). Similarly,

$$|\mathcal{E}[u_1^i(Y)] - \mathcal{E}[u_1^i(Z)]| \leq \|Y; Z\| \leq \frac{\epsilon}{8}. \quad (35)$$

Moreover, notice that the random variable Y satisfies the following condition

$$|\mathcal{E}[u_0^i(Y)] - \mathcal{E}[u_1^i(Y)]| \leq 3\epsilon/4, \quad (36)$$

since, in order for v_i to be included into \mathcal{S}_i , the test in Step 3(c)iv of MOMENT SEARCH must have succeeded. Combining (34), (35) and (36) we get (28). This concludes the proof. ■

Computational Complexity We will argue that there exists an implementation of MOMENT SEARCH which on input $k = O(1/\epsilon)$ runs in time

$$U \cdot \text{poly}(n) \cdot (1/\epsilon)^{O(\log^2(1/\epsilon))},$$

where U is the number of bits required to represent a payoff value of the game.

Observe first that the number of possible guesses for Step 1 of MOMENT SEARCH is at most $n^2 O((1/\epsilon)^6)$. Observe further that the number of possible guesses for μ_ℓ in Step 2 is at most

$t \left(\frac{k^2}{2}\right)^\ell$ (where $t \leq k^3$ is the number of players who mix), so jointly the number of possible guesses for μ_1, \dots, μ_d is at most

$$\prod_{\ell=1}^d t \left(\frac{k^2}{2}\right)^\ell = t^d \left(\frac{k^2}{2}\right)^{d(d+1)/2} = \left(\frac{1}{\epsilon}\right)^{O(\log^2 \frac{1}{\epsilon})}.$$

The same asymptotic upper bound applies to the total number of guesses for μ'_1, \dots, μ'_d . Given the above the total number of guesses that MOMENT SEARCH has to do is

$$n^2 \left(\frac{1}{\epsilon}\right)^{O(\log^2 \frac{1}{\epsilon})}.$$

We next argue that the running time required to complete Steps 3, 4, and 5 is at most

$$O(n^3) \cdot U \cdot \left(\frac{1}{\epsilon}\right)^{O(\log^2(1/\epsilon))}.$$

For this we establish the following; we give the proof in the end of this section.

Claim 11. *Given a set of values $\mu_1, \dots, \mu_d, \mu'_1, \dots, \mu'_d$, where, for all $\ell = 1, \dots, d$,*

$$\mu_\ell, \mu'_\ell \in \left\{0, \left(\frac{1}{k^2}\right)^\ell, 2\left(\frac{1}{k^2}\right)^\ell, \dots, B\right\},$$

for some $B \in \mathbb{N}$, discrete sets $\mathcal{T}_1, \dots, \mathcal{T}_m \subseteq \{0, \frac{1}{k^2}, \frac{2}{k^2}, \dots, 1\}$, and four integers $m_0, m_1 \leq m$, $m_s, m_b \leq B$, it is possible to solve the system of equations:

$$\begin{aligned} (\Sigma) : \quad & \sum_{p_i \in (0, 1/2]} p_i^\ell = \mu_\ell, \text{ for all } \ell = 1, \dots, d, \\ & \sum_{p_i \in (1/2, 1)} p_i^\ell = \mu'_\ell, \text{ for all } \ell = 1, \dots, d, \\ & |\{i | p_i = 0\}| = m_0 \\ & |\{i | p_i = 1\}| = m_1 \\ & |\{i | p_i \in (0, 1/2]\}| = m_s \\ & |\{i | p_i \in (1/2, 1)\}| = m_b \end{aligned}$$

with respect to the variables $p_1 \in \mathcal{T}_1, \dots, p_m \in \mathcal{T}_m$, or to determine that no solution exists, in time

$$O(m^3) B^{O(d)} k^{O(d^2)}.$$

Applying Claim 11 with $m \leq t$, $B \leq t$ (where $t \leq k^3$ is the number of players who mix), $m_0 = 0$, $m_1 = 0$, shows that Steps 3(a)i, 3(c)i can be completed in time

$$O(t^3) t^{O(d)} k^{O(d^2)} = \left(\frac{1}{\epsilon}\right)^{O(\log^2(1/\epsilon))}.$$

Another application of Claim 11 with $m = n$, $B \leq t$, $m_0 \leq n$, $m_1 \leq n$ shows that Step 4 of MOMENT SEARCH can be completed in time

$$O(n^3)t^{O(d)}k^{O(d^2)} = O(n^3) \cdot \left(\frac{1}{\epsilon}\right)^{O(\log^2(1/\epsilon))}.$$

Finally, we argue that the computation of the expected utilities \mathcal{U}_0^i and \mathcal{U}_1^i required in Steps 3(a)iii, 3(c)iii of MOMENT SEARCH can be done efficiently using dynamic programming with $O(n^2)$ operations on numbers with at most $b(n, k) := \lceil 1 + n \log_2(k^2) + U \rceil$ bits, where U is the number of bits required to specify a payoff value of the game.⁵

Therefore, the overall time required for the execution of MOMENT SEARCH is

$$O(n^3) \cdot U \cdot \left(\frac{1}{\epsilon}\right)^{O(\log^2(1/\epsilon))}.$$

Proof of Claim 11: We use dynamic programming. Let us consider the following tensor of dimension $2d + 5$:

$$A(i, z_0, z_1, z_s, z_b; \nu_1, \dots, \nu_d; \nu'_1, \dots, \nu'_d),$$

where $i \in [m]$, $z_0, z_1 \in \{0, \dots, m\}$, $z_s, z_b \in \{0, \dots, B\}$ and

$$\nu_\ell, \nu'_\ell \in \left\{0, \left(\frac{1}{k^2}\right)^\ell, 2\left(\frac{1}{k^2}\right)^\ell, \dots, B\right\},$$

for $\ell = 1, \dots, d$. The total number of cells in A is

$$\begin{aligned} m \cdot (m+1)^2 \cdot (B+1)^2 \cdot \left(\prod_{\ell=1}^d (Bk^{2\ell} + 1)\right)^2 \\ \leq O(m^3)B^{O(d)}k^{2d(d+1)}. \end{aligned}$$

Every cell of A is assigned value 0 or 1, as follows:

$$\begin{aligned} A(i, z_0, z_1, z_s, z_b; \nu_1, \dots, \nu_d, \nu'_1, \dots, \nu'_d) = 1 \\ \Leftrightarrow \left(\begin{array}{l} \text{There exist } p_1 \in \mathcal{T}_1, \dots, p_i \in \mathcal{T}_i \text{ such} \\ \text{that } |\{j \leq i | p_j = 0\}| = z_0, \\ \quad |\{j \leq i | p_j = 1\}| = z_1, \\ \quad |\{j \leq i | p_j \in (0, 1/2]\}| = z_s, \\ \quad |\{j \leq i | p_j \in (1/2, 1)\}| = z_b, \\ \quad \sum_{j \leq i: p_j \in (0, 1/2]} p_j^\ell = \nu_\ell, \text{ for all} \\ \ell = 1, \dots, d, \sum_{j \leq i: p_j \in (1/2, 1)} p_j^\ell = \nu'_\ell, \text{ for} \\ \text{all } \ell = 1, \dots, d. \end{array} \right). \end{aligned}$$

⁵To compute a bound on the number of bits required for the expected utility computations, note that every non-zero probability value that is computed along the execution of the algorithm must be an integer multiple of $(\frac{1}{k^2})^{n-1}$, since the mixed strategies of all players are from the set $\{0, 1/k^2, 2/k^2, \dots, 1\}$. Further note that the expected utility is a weighted sum of $(n-1)$ payoff values, with U bits required to represent each value, and all weights being probabilities.

It is easy to complete A working in layers of increasing i . We initialize all entries to value 0. Then, the first layer $A(1, \cdot, \cdot; \cdot, \dots, \cdot)$ can be completed easily as follows:

$$\begin{aligned} A(1, 1, 0, 0, 0; 0, 0, \dots, 0; 0, 0, \dots, 0) &= 1 \Leftrightarrow 0 \in \mathcal{T}_1 \\ A(1, 0, 1, 0, 0; 0, 0, \dots, 0; 0, 0, \dots, 0) &= 1 \Leftrightarrow 1 \in \mathcal{T}_1 \\ A(1, 0, 0, 1, 0; p, p^2, \dots, p^d; 0, \dots, 0) &= 1 \Leftrightarrow p \in \mathcal{T}_1 \cap (0, 1/2] \\ A(1, 0, 0, 0, 1; 0, \dots, 0; p, p^2, \dots, p^d) &= 1 \Leftrightarrow p \in \mathcal{T}_1 \cap (1/2, 1) \end{aligned}$$

Inductively, to complete layer $i + 1$, we consider all the non-zero entries of layer i and for every such non-zero entry and for every $v_{i+1} \in \mathcal{T}_{i+1}$, we find which entry of layer $i + 1$ we would transition to if we chose $p_{i+1} = v_{i+1}$. We set that entry equal to 1 and we also save a pointer to this entry from the corresponding entry of layer i , labeling that pointer with the value v_{i+1} . The time we need to complete layer $i + 1$ is bounded by

$$|\mathcal{T}_{i+1}| \cdot (m + 1)^2 B^{O(d)} k^{2d(d+1)} \leq O(m^2) B^{O(d)} k^{O(d^2)}.$$

Therefore, the overall time needed to complete A is

$$O(m^3) B^{O(d)} k^{O(d^2)}.$$

After completing tensor A , it is easy to check if there exists a solution to (Σ) . A solution exists if and only if

$$A(m, m_0, m_1, m_s, m_b; \mu_1, \dots, \mu_d; \mu'_1, \dots, \mu'_d) = 1,$$

and can be found by tracing back the pointers from this cell of A . The overall running time is dominated by the time needed to fill in A . ■